SERIES G: TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

Packet over Transport aspects – Synchronization, quality and availability targets

SERIES Y: GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS, NEXT-GENERATION NETWORKS, INTERNET OF THINGS AND SMART CITIES

Internet protocol aspects – Transport

Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network

Amendment 2

ITU-T G-SERIES RECOMMENDATIONS

TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

INTERNATIONAL TELEPHONE CONNECTIONS AND CIRCUITS	G.100–G.199
GENERAL CHARACTERISTICS COMMON TO ALL ANALOGUE CARRIER-TRANSMISSION SYSTEMS	G.200–G.299
INDIVIDUAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON METALLIC LINES	G.300–G.399
GENERAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON RADIO-RELAY OR SATELLITE LINKS AND INTERCONNECTION WITH METALLIC LINES	G.400–G.449
COORDINATION OF RADIOTELEPHONY AND LINE TELEPHONY	G.450–G.499
TRANSMISSION MEDIA AND OPTICAL SYSTEMS CHARACTERISTICS	G.600–G.699
DIGITAL TERMINAL EQUIPMENTS	G.700–G.799
DIGITAL NETWORKS	G.800–G.899
DIGITAL SECTIONS AND DIGITAL LINE SYSTEM	G.900–G.999
MULTIMEDIA QUALITY OF SERVICE AND PERFORMANCE – GENERIC AND USER-RELATED ASPECTS	G.1000–G.1999
TRANSMISSION MEDIA CHARACTERISTICS	G.6000–G.6999
DATA OVER TRANSPORT – GENERIC ASPECTS	G.7000–G.7999
PACKET OVER TRANSPORT ASPECTS	G.8000–G.8999
Ethernet over Transport aspects	G.8000–G.8099
MPLS over Transport aspects	G.8100–G.8199
Synchronization, quality and availability targets	**G.8200–G.8299**
Mobile network transport aspects	G.8300–G.8399
Service Management	G.8600–G.8699
ACCESS NETWORKS	G.9000–G.9999

For further details, please refer to the list of ITU-T Recommendations.

Recommendation ITU-T G.8275.2/Y.1369.2

Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network

Amendment 2

Summary

Amendment 2 to this Recommendation provides the following updates:

–   Details added on the use of [IEEE 1588-2019]

–   New profileVersion 1.2 and new profile Identifier added for [IEEE 1588-2019]

–   New appendix added "Considerations on selecting time out values"

–   New material added to clause 6.7.11 *Packet timing signal fail*.

–   New member "defaultDS.sdoId" added to Table A.1 *defaultDS data set member specifications*

–   New member "currentDS.meanDelay" added to Table A.2 *currentDS data set member specifications*

–   New member "portDS.minorVersionNumber" added to Table A.5 *portDS data set member specifications*

–   Parameters for "portDS.syncReceiptTimeout" and "portDS.delayRespReceiptTimeout" in Table A.5, which were "for further study", were filled in.

History

Edition	Recommendation	Approval	Study Group	Unique ID[*]
1.0	ITU-T G.8275.2/Y.1369.2	2016-06-22	15	11.1002/1000/12833
1.1	ITU-T G.8275.2/Y.1369.2 (2016) Amd. 1	2017-08-29	15	11.1002/1000/13330
1.2	ITU-T G.8275.2/Y.1369.2 (2016) Amd. 2	2018-03-16	15	11.1002/1000/13556
1.3	ITU-T G.8275.2/Y.1369.2 (2016) Amd. 3	2019-08-29	15	11.1002/1000/14018
2.0	ITU-T G.8275.2/Y.1369.2	2020-03-15	15	11.1002/1000/14216
2.1	ITU-T G.8275.2/Y.1369.2 (2020) Amd. 1	2020-11-13	15	11.1002/1000/14544
2.2	ITU-T G.8275.2/Y.1369.2 (2020) Amd. 2	2021-06-29	15	11.1002/1000/14710

Keywords

IEEE 1588, partial timing support, phase and time synchronization, PTP, telecom profile.

---

[*]   To access the Recommendation, type the URL http://handle.itu.int/ in the address field of your web browser, followed by the Recommendation's unique ID. For example, http://handle.itu.int/11.1002/1000/11830-en.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure, e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had received notice of intellectual property, protected by patents/software copyrights, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the appropriate ITU-T databases available via the ITU-T website at http://www.itu.int/ITU-T/ipr/.

Table of Contents

Recommendation ITU-T G.8275.2/Y.1369.2

Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network

Amendment 2

Editorial note: This is a complete-text publication. Modifications introduced by this amendment are shown in revision marks relative to Recommendation ITU-T G.8275.2/Y.1369.2 (2020) plus its Amendment 1.

1   Scope

This Recommendation specifies a profile for telecommunication applications based on [IEEE 1588] precision time protocol (PTP). The profile specifies the IEEE 1588 functions that are necessary to ensure network element interoperability for the delivery of accurate phase/time (and frequency) synchronization. The profile is based on the use of partial timing support (PTS) or assisted partial timing support (APTS) from the network architecture as described in [ITU-T G.8275] and definitions described in [ITU-T G.8260].

It is assumed that this profile will be used in well-planned cases where network behaviour and performance can be constrained within well-defined limits, including limits on static asymmetry. Control of static asymmetries can be achieved in case of assisted partial timing support. Use of this profile in unassisted mode would require careful considerations on how to control static asymmetries. Additional considerations are included in [ITU-T G.8271.2]. This version of the profile specifies the high-level design requirements, modes of operation for the exchange of PTP messages, the PTP protocol mapping, the best master clock algorithm (BMCA) options, as well as the PTP protocol configuration parameters.

At the time of publication of this profile, performance analysis, network limits, and clocks used in the profile, namely boundary and slave clocks, are for further study.

This Recommendation also specifies some aspects necessary for use in a telecom environment that are outside the scope of the PTP profile but complement it.

NOTE – [b-IEEE 1588-2019] specifies a new PTP version 2.1. Based on clause 19 of [b-IEEE 1588-2019], PTP version 2.1 is inter-operable with PTP version 2.0, which is specified in IEEE Std 1588-2008 [IEEE 1588]. Therefore, it is expected that PTP version 2.1 can be used in this profile. Further details about the use of PTP version 2.1 in this Recommendation are for further study. The intention is that an ITU-T G.8275.2 implementation based on PTP version 2.1 will be compatible/inter-operable with existing deployments using PTP version 2.0, as long as no new PTP functionality that is not present in PTP version 2.0 is used.

An implementation compliant with this profile can claim compliance with either IEEE Std 1588-2008 [IEEE 1588-2008] or IEEE Std 1588-2019 [IEEE 1588-2019]. Considerations on the use of one or the other profile are provided in Appendix IX. Compliance for a specific implementation with either the [IEEE1588-2008] or [IEEE1588-2019] versions of the standard should be stated when referring to this profile.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

[ITU-T G.781]	Recommendation ITU-T G.781 (2017), *Synchronization layer functions*.
[ITU-T G.810]	Recommendation ITU-T G.810 (1996), *Definitions and terminology for synchronization networks*.
[ITU-T G.8260]	Recommendation ITU-T G.8260 (2020), *Definitions and terminology for synchronization in packet networks*.
[ITU-T G.8265.1]	Recommendation ITU-T G.8265.1/Y.1365.1 (2014), *Precision time protocol telecom profile for frequency synchronization*.
[ITU-T G.8271]	Recommendation ITU-T G.8271/Y.1366 (2020), *Time and phase synchronization aspects of telecommunication networks*.
[ITU-T G.8271.2]	Recommendation ITU-T G.8271.2/Y.1366.2 (2017), *Network limits for time synchronization in packet networks with partial timing support from the network*.
[ITU-T G.8272]	Recommendation ITU-T G.8272/Y.1367 (2018), *Timing characteristics of primary reference time clocks*.
[ITU-T G.8272.1]	Recommendation ITU-T G.8272.1/Y.1367.1 (2016), *Timing characteristics of enhanced primary reference time clocks*.
[ITU-T G.8273]	Recommendation ITU-T G.8273/Y.1368 (2013), *Framework of phase and time clocks*.
[ITU-T G.8273.4]	Recommendation ITU-T G.8273.4/Y.1368.4 (2020), *Timing characteristics of telecom boundary clocks and telecom time slave clocks for use with partial timing support from the network*.
[ITU-T G.8275]	Recommendation ITU-T G.8275/Y.1369 (2020), *Architecture and requirements for packet-based time and phase distribution*.
[ITU-T G.8275.1]	Recommendation ITU-T G.8275.1/Y.1369.1 (2020), *Precision time protocol telecom profile for phase/time synchronization with full timing support from the network*.
[IEEE 1588]	Either [IEEE 1588-2008] or [IEEE 1588-2019] depending on the specific implementation. See clause 5 Conventions for more details.
[IEEE 1588-2008]	IEEE 1588-2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*.
[IEEE 1588-2019]	IEEE 1588-2019, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*.

3 Definitions

3.1 Terms defined elsewhere

The terms and definitions used in this Recommendation are contained in [ITU-T G.810] and [ITU-T G.8260].

3.2 Terms defined in this Recommendation

None.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

APTS	Assisted Partial Timing Support
BC	Boundary Clock
BMCA	Best Master Clock Algorithm
ePRTC	Enhanced Primary Reference Time Clock
EUI	Extended Unique Identifier
GM	GrandMaster
GNSS	Global Navigation Satellite System
IP	Internet Protocol
MA-L	MAC Address – Large
MA-M	MAC Address – Medium
MA-S	MAC Address – Small
OC	Ordinary Clock
OUI	Organizationally Unique Identifier
ParentDS	Parent Data Set
PDV	Packet Delay Variation
PRC	Primary Reference Clock
PRS	Primary Reference Source
PRTC	Primary Reference Time Clock
PTP	Precision Time Protocol
PTPVAR	PTP Variance
PTS	Partial Timing Support
PTSF	Packet Timing Signal Fail
QL	Quality Level
SDH	Synchronous Digital Hierarchy
SF	Signal Fail
SSM	Synchronization Status Message
SSU	Synchronization Supply Unit
SSU-A	Primary level SSU
SSU-B	Secondary level SSU
ST2	Stratum 2
ST3E	Stratum 3 Enhanced
T-BC-A	Telecom Boundary Clock for Assisted partial timing support

T-BC-P   Telecom Boundary Clock for Partial timing support

TC   Transparent Clock

T-GM   Telecom Grandmaster

TLV   Type, Length, Value

T-TC-P   Telecom Transparent Clock for Partial timing support

T-TSC-A   Telecom Time Slave Clock for Assisted partial timing support

T-TSC-P   Telecom Time Slave Clock for Partial timing support

UDP   User Datagram Protocol

VLAN   Virtual Local Area Network

5   Conventions

Within this Recommendation, the following conventions are used: the term PTP refers to the PTP ~~version 2~~ protocol defined in [IEEE 1588]. PTP messages used within this Recommendation are defined in [IEEE 1588] and are identified using italicized text.

The term telecom grandmaster (T-GM) refers to a device consisting of a grandmaster (GM) clock as defined in [IEEE 1588] and this Recommendation, with additional performance characteristics for further study.

The term telecom boundary clock for partial timing support (T-BC-P) refers to a device consisting of a boundary clock (BC) as defined in [IEEE 1588], with additional performance characteristics as defined in [ITU-T G.8273.4]. The term telecom boundary clock for assisted partial timing support (T-BC-A) refers to a device consisting of a boundary clock (BC) as defined in [IEEE 1588] that is assisted by a local time reference (e.g., a PRTC or GNSS-based time source) as a primary source of time, with additional performance characteristics as defined in [ITU-T G.8273.4].

NOTE 1 – In comparing the operation of the T-BC-P and the T-BC-A, the T-BC-P uses PTP input as the primary source of synchronization. The T-BC-A is targeted for the APTS scenario described in [ITU-T G.8271.2], where the PTP input is used only as a secondary source of synchronization to hold the time for up to 72 hours, and is not intended to use PTP as the primary timing source.

The term telecom transparent clock for partial timing support (T-TC-P) refers to a device consisting of a transparent clock (TC) as defined in [IEEE 1588], with additional performance characteristics for further study.

The term telecom time slave clock for partial timing support (T-TSC-P) refers to a device consisting of either an ordinary clock (OC), with one PTP port, or a boundary clock (BC), with multiple PTP ports, as defined in [IEEE 1588] and this Recommendation, that does not support providing synchronization using PTP to other PTP clocks in the PTP domain, and with additional performance characteristics as defined in [ITU-T G.8273.4].

The term telecom time slave clock for assisted partial timing support (T-TSC-A) refers to a T-TSC-P that is assisted by a local time reference (e.g., a PRTC or GNSS-based time source) as a primary source of time, with additional performance characteristics as defined in [ITU-T G.8273.4]. Note that in the case of a T-TSC-A or T-TSC-P with multiple PTP ports (BC), only one PTP port can be in PTP SLAVE state at any instant in time based on the BMCA. Other PTP ports not in the PTP SLAVE state may actively exchange synchronization messages with other PTP clocks populated in the unicast master table using unicast negotiation.

NOTE 2 – In comparing the operation of the T-TSC-P and the T-TSC-A, the T-TSC-P uses the PTP input as the primary source of synchronization. The T-TSC-A is targeted for the APTS scenario described in [ITU-T G.8271.2], where the PTP input is used only as a secondary source of synchronization to hold the time for up to 72 hours, and is not intended to use PTP as the primary timing source.

The term primary reference time clock (PRTC) refers to the clock defined in [ITU-T G.8272]. The term enhanced primary reference time clock (ePRTC) refers to an enhanced version of the PRTC, which is being studied.

In this Recommendation, T-BC-P/A means T-BC-P or T-BC-A. Likewise, T-TSC-P/A means T-TSC-P or T-TSC-A.

Within this recommendation, some requirements are stated as requiring compliance to [IEEE 1588]. For implementations compliant to [IEEE 1588-2008], the reference to [IEEE 1588] means compliance to [IEEE 1588-2008]. For implementations compliant to [IEEE 1588-2019], the reference to [IEEE 1588] means compliance to [IEEE 1588-2019]. Some of these references to [IEEE 1588] include a specific clause number. In these cases, the clause number is the same in both [IEEE 1588-2008] and [IEEE 1588-2019]. If the requirements are in different clauses in the two versions of IEEE 1588, then the text of this Recommendation shall include the specific clause for [IEEE 1588-2008] and the specific cause for [IEEE 1588-2019].

6   Use of PTP for phase/time distribution

The 2002 version of the [IEEE 1588] standard defines was developed by the IEEE initially to support the timing requirements of industrial automation and test and measurement, defining the precision time protocol (PTP) designed to enable accurate time transfer in this context.

  IEEE 1588-2008 (defined in [IEEE 1588]) contains features useful to the transport of the protocol over a wide area network, and It introduces the concept of "profile", whereby aspects of the protocol may be selected and specified for a particular use other than the originally intended industrial automation application.

A PTP profile was defined by ITU-T in [ITU-T G.8265.1] to address applications requiring frequency synchronization only. An additional PTP profile was defined by ITU-T in [ITU-T G.8275.1] in order to allow the distribution of phase/time with full timing support from the network. This Recommendation defines another PTP profile to allow the distribution of phase and time with partial timing support (PTS) from the network.

The [IEEE 1588] telecom profile defined within this Recommendation is intended to be used by telecom applications requiring accurate phase and time synchronization. It covers applications where there is need for phase alignment and/or time of day. It supports the specific architecture described in [ITU-T G.8275] in order to allow the distribution of phase/time with PTS from the network and is based on the 2008 version of PTP defined in [IEEE 1588]. This includes the case of assisted partial timing support (APTS).

This profile uses only the unicast mode.

In order to claim compliance with the telecom profile, the requirements of this Recommendation and the relevant requirements of [IEEE 1588], as referenced in Annex A, must be met.

The detailed aspects related to the telecom profile are described in the following clauses, while the profile itself is contained in Annex A. It follows the general rules for profile specification developed in [IEEE 1588].

This PTP telecom profile defines the [IEEE 1588] parameters to be used, in order to guarantee protocol interoperability between implementations and specifies the optional features, default values of configurable attributes and mechanisms that must be supported. However, it does not guarantee that the performance requirements of a given application will be met. Those performance aspects are currently under study and imply additional elements beyond the content of the PTP profile itself. These are planned to be addressed in other ITU-T Recommendations.

6.1 High-level design requirements

~~Clause 19.3.1.1 of~~ [IEEE 1588] states:

"The purpose of a PTP profile is to allow organizations to specify specific selections of attribute values and optional features of PTP that, when using the same transport protocol, inter-work and achieve a performance that meets the requirements of a particular application."

For operation in a telecom network, some additional criteria are also required to be consistent with standard telecom synchronization practices. With that in mind, the PTP profile for time and phase distribution must meet the following high-level requirements:

1) Mechanisms must be specified to allow interoperability between the various phase/time clocks belonging to the architecture defined in [ITU-T G.8275] and described in [ITU-T G.8273].

2) Mechanisms must permit consistent operation over managed wide area telecom networks.

3) Packet-based mechanisms must allow the synchronization network to be designed and configured in a fixed arrangement.

4) Protection schemes used by packet-based systems must be based on standard telecom operational practice and allow T-TSC-P and T-TSC-A to have the ability to take phase and time from multiple geographically separate T-GM clocks.

5) Phase/time reference source selection based on received phase/time traceability and local priority, as well as automatic establishment of the phase/time synchronization network topology, should be permitted.

6.2 PTP modes and options

6.2.1 PTP Domains

A domain consists of a logical grouping of clocks communicating with each other using the PTP protocol.

PTP domains are used to partition a network within an administrative domain. The PTP messages and data sets are associated with a domain and therefore the PTP protocol is independent for different domains.

In this PTP telecom profile, the default PTP domain number is 44, and the range of applicable PTP domain numbers is {44 – 63}.

NOTE – This range has been selected from the user-defined PTP domain number range defined in [IEEE 1588]. Although non-overlapping ranges have been considered for the different PTP telecom profiles so that interactions between the profiles are prevented, nothing precludes another industry from using the same user-defined PTP domain number range when defining a non-telecom PTP profile. It is the responsibility of the network operator to identify if the risk of unintentional interactions between PTP profiles exists, and to take the necessary actions to prevent such behaviour.

6.2.2 PTP messages

[IEEE 1588] defines two categories of message types: event and general PTP messages. The two types differ in that event messages are timed messages and require or contain an accurate timestamp. General message types do not require accurate timestamps.

[IEEE 1588] defines the following message types: *Sync*, *Delay_Req* (i.e., "delay request"), *Announce*, *Follow_Up*, *Delay_Resp* (i.e., "delay response"), *Pdelay_Req*, *Pdelay_Resp*, and *Pdelay_Resp_Follow_Up*, *Management* and *Signalling*.

Sync, Delay_Req, Announce, Follow_Up, Delay_Resp, and Signalling messages are used in this profile.

Pdelay_Req, *Pdelay_Resp*, and *Pdelay_Resp_Follow_Up* messages are not used in this profile.

The use of *Management* messages is for further study.

6.2.3 Types of PTP clocks supported in the profile

The OC and BC according to [IEEE 1588] are used in this profile.

There are two types of OCs:

1) OC that can only be a grandmaster (T-GM according to the architecture defined in [ITU-T G.8275], and as included in [ITU-T G.8272]).

2) OC that can only be a slave, i.e., slave-only OC (T-TSC-P with only one port or T-TSC-A with only one port according to the architecture defined in [ITU-T G.8275]). The clock specifications for T-TSC-P and T-TSC-A are defined in [ITU-T G.8273.4].

There are three types of BCs:

1) BC that can only be a grandmaster (T-GM according to the architecture defined in [ITU-T G.8275], and as included in [ITU-T G.8272]).

2) BC that can become a grandmaster and can also be slaved to another PTP clock (T-BC-P and T-BC-A according to architecture defined in [ITU-T G.8275]). The clock specifications for the T-BC-P and T-BC-A are defined in [ITU-T G.8273.4].

3) BC that can only be a slave (T-TSC-P with more than one port or T-TSC-A with more than one port according to the architecture defined in [ITU-T G.8275]). The clock specifications for T-TSC-P and T-TSC-A are defined in [ITU-T G.8273.4].

NOTE – T-GM and GM are different concepts; GM is a status defined in [IEEE 1588] that a PTP clock may obtain if it wins the BMCA, while T-GM is a type of clock defined in the [ITU-T G.8275] architecture.

The mapping between these PTP clock types and the phase/time clocks defined in the [ITU-T G.8275] architecture is described in Table 1.

Table 1 − Mapping between ITU-T G.8275.2 and PTP clock types

Clock type from [ITU-T G.8275.2]	Description	Clock type from [IEEE 1588]
T-GM	Master-only ordinary clock (master with a single PTP port, cannot be slaved to another PTP clock)	OC
	Master-only boundary clock (master with multiple PTP ports, cannot be slaved to another PTP clock)	BC (Note 1)
T-BC-P (partial)	Boundary clock (may become a GM, or may be slaved to another PTP clock)	BC
T-BC-A (assisted partial)	Boundary clock assisted by a local time reference that is used as a primary source of time (may become a GM, or may be slaved to another PTP clock)	BC (Note 2)
T-TSC-P (partial)	Slave-only, single port, ordinary clock (always a slave)	OC
	PTP clock at the end of the PTP synchronization chain, multiple port clock	BC (Note 1)
T-TSC-A (assisted partial)	Slave-only, single port, ordinary clock (always a slave) assisted by a local time reference that is used as a primary source of time	OC (Note 2)

Table 1 – Mapping between ITU-T G.8275.2 and PTP clock types

Clock type from [ITU-T G.8275.2]	Description	Clock type from [IEEE 1588]
	PTP clock at the end of the PTP synchronization chain, multiple port clock assisted by a local time reference that is used as a primary source of time	BC (Note 1) (Note 2)
NOTE 1 – According to [IEEE 1588], a clock that has multiple PTP ports is by definition a boundary clock. NOTE 2 – Examples of a local time reference: PRTC or GNSS-based time source.		

6.3 PTP modes

[IEEE 1588] describes several modes of operation between a master-port (which is a PTP in MASTER state) and a slave-port (which is a PTP port in SLAVE state). The term grant-port refers to a PTP port granting and providing PTP message service, and the term request-port refers to a PTP port requesting and receiving PTP message service. Typically, the grant-port is a master-port and the request-port is a slave-port. Information related to grant-ports and request-ports in other PTP states will be included in a future version of this Recommendation related to PTP clocks with multiple PTP ports.

NOTE 1 – A grant-port may be in the MASTER state, PASSIVE state, LISTENING state, PRE_MASTER state, UNCALIBRATED state, or SLAVE state. (but not INITIALIZING, FAULTY, or DISABLED state).

NOTE 2 – A request-port may be in the MASTER state, PASSIVE state, LISTENING state, PRE_MASTER state, UNCALIBRATED state, or SLAVE state. (but not INITIALIZING, FAULTY, or DISABLED state).

This clause describes these modes with respect to functionality needed to be compliant with this profile.

6.3.1 One-way versus two-way operation

A PTP master-port or grant-port compliant with the profile must be capable of supporting one-way and two-way timing transfers. For APTS, since only PTP synchronization may be required, a slave-port or request-port may only utilize one-way mode, or may utilize two-way mode, but is not required to support both methods; otherwise for PTS, a slave-port or request-port must utilize two-way.

NOTE – In the APTS case, even if performance objectives are specified by means of two-way metrics, this does not prevent the slave-port or request-port from utilizing one-way mode, although for a more accurate interpretation of how the network characteristics relates to the expected performance of the clock, two-way operation may be preferred.

6.3.2 One-step versus two-step clock mode

PTP defines two types of clock behaviour: the "one-step clock" and the "two-step clock". In a one-step clock, the precise timestamp is transported directly in the *Sync* message. In a two-step clock, a *Follow_Up* message is used to carry the precise timestamp of the corresponding *Sync* message. The use of *Follow_Up* messages is optional in the PTP protocol.

The one-step clock approach enables equipment to reduce significantly the number of PTP messages sent by the master-port or grant-port and relax the master-port or grant-port capacities.

However, there might be situations where the two-step clock approach might be required (e.g., when some security features are required). These situations are for further study.

Both one-step and two-step clocks are allowed in the profile. A PTP master-port or grant-port compliant with the profile may use either a one-step clock or a two-step clock or both.

NOTE – The performance of the PTP timing flow generated by the master-port or grant-port with those two approaches is for further study.

To be compliant with [IEEE 1588], a slave-port or request-port must be capable of handling both one-step clock and two-step clock, without any particular configuration.

As per ~~clause 7.3.8.3 of~~ [IEEE 1588], when a two-step clock is used, the value of the flag "twoStepFlag" shall be TRUE to indicate that a *Follow_up* message will follow the *Sync* message, and that the slave-port or request-port must not consider the originTimestamp embedded in the *Sync* message. When a one-step clock is used, the value of the flag "twoStepFlag" shall be FALSE, and the slave-port or request-port must consider the originTimestamp embedded in the *Sync* message in this case.

6.3.3 Unicast versus multicast mode

PTP allows the use of unicast and multicast modes for the transmission of the PTP messages.

For the PTP profile specified in Annex A, the unicast mode is used for all the PTP messages.

A master-port or grant-port compliant with the PTP profile specified in Annex A must support the unicast mode.

A slave-port or request-port compliant with the PTP profile specified in Annex A must support the unicast mode.

6.4 PTP mapping

This PTP telecom profile is based on the PTP mapping ~~defined in [IEEE 1588] Annex D,~~ *Transport of PTP over User Datagram Protocol over Internet Protocol Version 4* [IEEE 1588] and ~~[IEEE 1588] Annex E,~~ *Transport of PTP over User Datagram Protocol over Internet Protocol Version 6* [IEEE 1588].

~~Therefore, a~~ A master-port, grant-port, slave-port, or a request-port compliant with the profile described in this Recommendation must be compliant with *Transport of PTP over User Datagram Protocol over Internet Protocol Version 4* [IEEE 1588] ~~[IEEE 1588] Annex D~~ and may be compliant with *Transport of PTP over User Datagram Protocol over Internet Protocol Version 6* [IEEE 1588]~~[IEEE 1588] Annex E~~.

NOTE – The use of the Internet Protocol (IP)/user datagram protocol (UDP) mapping is to facilitate the use of IP addressing. It does not imply that the PTP flow can be carried over an unmanaged packet network. It is assumed that a well-controlled packet network will be used to control and minimize packet delay variation.

6.5 Message rates

The message rate values are only defined for protocol interoperability purposes. It is not expected that any slave clock shall meet the relevant target performance requirements at all packet rates within the given range, specifically at the lower packet rate. The appropriate value depends on the clock characteristics and on the target performance requirements. Different packet rate needs may also apply during the stabilization period.

NOTE – A specific slave clock implementation, in order to meet its target performance requirements, may support a subset of the message rates within the ranges noted below. A master-port or grant-port, on the other hand, is required to support the full range of message transmission rates. Unless an implementation specifies otherwise, the default value listed below is assumed to be used.

Within the scope of the profile, the following messages can be used and the corresponding indicated range of rates must be respected for unicast messages:

– *Sync* messages (if used, *Follow_up* messages will have the same rate) – minimum rate: 1 packet-per-second, maximum rate: 128 packets-per-second.

– *Delay_Req*/*Delay_Resp* messages – minimum rate:1 packet-per-second, maximum rate: 128 packets-per-second.

– *Announce* messages – minimum rate:1 packet-per-second, maximum rate: eight packets-per-second.

– *Signalling* messages – no rate is specified.

The use of *Management* messages is for further study.

6.6   Unicast message negotiation

Within a telecommunication network, there are benefits to allowing PTP request-ports to request the synchronization service from PTP grant-ports. [IEEE 1588] ~~offers~~ defines a unicast message negotiation mechanism to allow request-ports to request this service within a unicast environment ~~(see [IEEE 1588] clause 16.1)~~. This profile supports the unicast message negotiation in accordance with [IEEE 1588] and as described below.

PTP clocks compliant with the profile must support the unicast message negotiation mechanism as per clause 16.1 of [IEEE 1588] and as described in this clause.

When using the unicast mode, PTP request-ports request synchronization service by sending a PTP *Signalling* message in unicast, containing the REQUEST_UNICAST_TRANSMISSION type, length, value (TLV), to the IP address of the selected PTP grant-port.

NOTE 1 – In this telecom profile, unicast connection establishment without negotiation is for further study.

The *Signalling* message containing the REQUEST_UNICAST_TRANSMISSION TLV is periodically renewed.

When initiating unicast negotiation with a grant-port, a request-port can use all 1's as the initial value for the targetPortIdentity field of the Signalling message. Based on the response from the grant-port, the request-port can then learn the clockIdentity and portNumber of the grant-port and may use this in any subsequent Signalling message. The request-port may also continue to use all 1's. Similarly, the grant-port may either learn and use the clockIdentity and portNumber of the request-port, or use all 1's value for the targetPortIdentity field of the Signalling messages that it sends. Both grant-port and request-port must be prepared to handle both situations in reception, i.e., receive PTP Signalling messages with either their own clockIdentity and portNumber or with all 1's values for the targetPortIdentity field. A request-port should stop using a previously learned clockIdentity and portNumber when the established unicast session either expires or is cancelled, and the request-port has exhausted the re-try process described in clause 6.6. As indicated above, when a unicast session is granted, a grant-port may use (a) all 1's, or (b) the clockIdentity and portNumber from the sourcePortIdentity field of the received unicast session request, as the targetPortIdentity in the response.

The logInterMessagePeriod can be configured to adjust the requested transmission rate of *Sync*, *Announce* and *Delay_Resp* messages.

The configurable range for the logInterMessagePeriod is given in Annex A for all the relevant messages.

The durationField value in each REQUEST_UNICAST_TRANSMISSION TLV has a default initialization value of 300 seconds and a configurable range of 60 to 1000 seconds.

In the event that a PTP grant-port is unable to meet a given request-port request, it should deny the request entirely rather than offer the request-port less than it originally requested.

In the event of being denied service by a grant-port, or receiving no response to the service request:

– A request-port should wait a minimum of one second (after denial or no response received) before issuing a new unicast service request for that message type to the same grant-port.

– If a request-port has issued three service requests for the same message type with a "grant denied" response, it should either:

- cancel any granted unicast service it may have for other message types, and request service from a different grant-port, or
- wait a further 60 seconds before re-issuing the request to the same grant-port.

An example of the message exchange to initiate the unicast synchronization service is shown in Figure 1. The timing diagram example represents the exchange of unicast messages for a one-step clock (i.e., no *Follow_up* messages) using one-way mode (i.e., no *Delay_Req* or *Delay_Resp*).

The example shows a unicast negotiation process for a packet request-port sending *Signalling* messages for *Announce* and *Sync* requests; a packet grant-port granting the packet request-port the requested message rates; a packet grant-port transmitting the requested *Announce* and *Sync* message rates and the renewal of *Announce* and *Sync* before the expiration of durationField.

Note that several timing diagrams could be represented based on various exchanges of message types, the use of single or concatenated TLVs in *Signalling* messages, the use of different durationFields for each message type, etc. Figure 1 provides an example of message interaction; it is for illustrative purposes only and does not represent a particular implementation.
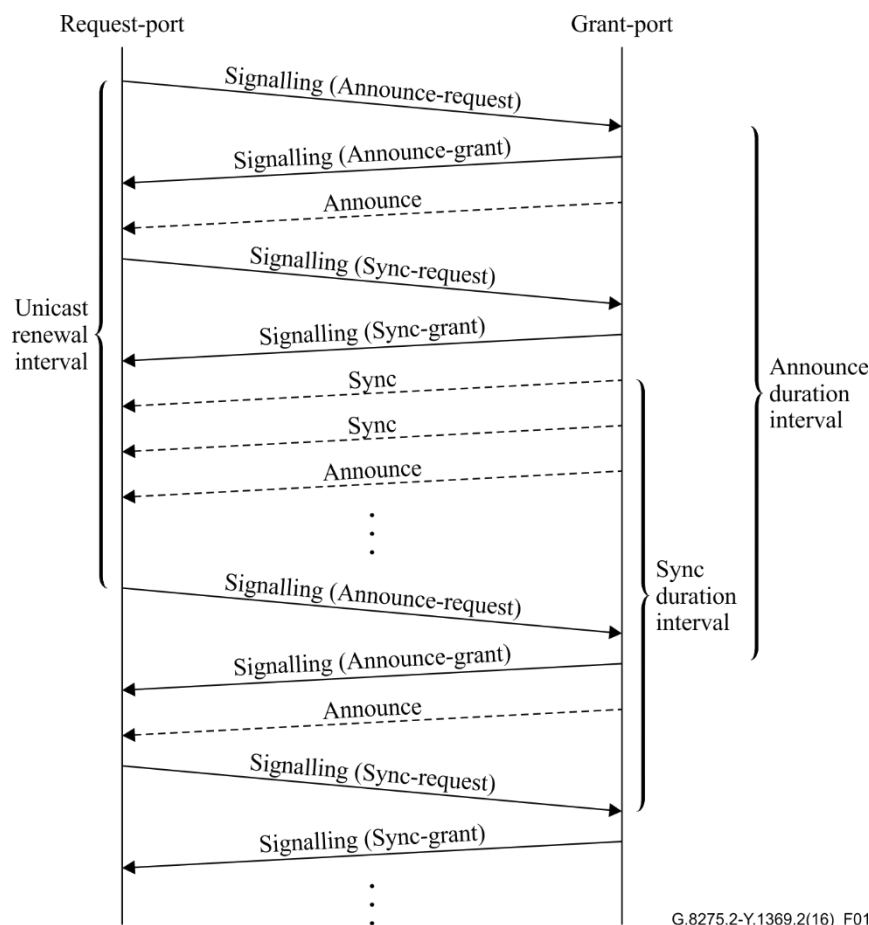
Figure 1 – Unicast negotiation example

PTP request-ports may request several types of PTP messages from a PTP grant-port (e.g., request-port working in two-way mode, which may request *Sync* and *Delay_Resp* messages, or request-port requesting *Announce* and *Sync* messages from the same grant-port). To request unicast transmission of different PTP message types, and to respond to such requests, [IEEE 1588] allows the use of a single *Signalling* message containing multiple TLVs or the use of multiple *Signalling* messages. Grant-ports and request-ports compliant with this profile must be prepared to handle those two situations. The expected behaviour during the initial negotiation and during the consecutive unicast service renewals is described in the paragraphs that follow.

Each request for unicast transmission from a specific request-port to a grant-port should start by issuing an *Announce* service type request first for that specific grant-port. Only after the request-port has been granted unicast service for the *Announce* message and received the first unicast *Announce* message from the specified grant-port, can the rest of the service type request take place. Such practice would ensure that the attributes (e.g., clockQuality) and capabilities of the specified grant-port are acceptable from the request-port's perspective before the rest of the services are contracted.

Upon receiving the first *Announce* message from the grant-port, the first *Signalling* message containing a REQUEST_UNICAST_TRANSMISSION TLV issued by the request-port should include all the service types the specific request-port requires from the grant-port using multiple REQUEST_UNICAST_TRANSMISSION TLVs. Such practice will reduce the chance that the grant-port will only grant part of the requested services in case it has been over-subscribed (due to simultaneous requests from other request-ports). The grant-port is allowed to respond to this request either with a single *Signalling* message containing multiple TLVs, or with multiple *Signalling* messages (e.g., each containing a single TLV).

When renewing the unicast services, the request-port, in sending *Signalling* messages (for 'keep-alive' purposes), may either continue to request all service types with a single *Signalling* message containing multiple TLVs, or with multiple independent *Signalling* messages (e.g., each containing a single TLV). The grant-port is allowed to respond to requests either with a single *Signalling* message containing multiple TLVs, or with multiple *Signalling* messages (e.g., each containing a single TLV).

~~The following text provided in clause A.9.4.2 of [IEEE 1588] should be followed: "For receiving continuous service, a requester should reissue a request in advance of the end of the grant period. The recommended advance should include sufficient margin for reissuing the request at least two more times if no grant is received."~~

As defined in [IEEE 1588], in order to receive continuous service, a requester should reissue a request in advance of the end of the grant period. The recommended advance should include sufficient margin for reissuing the request at least two more times if no grant is received.

In case the unicast transmission sessions are cancelled as defined in ~~clause 16.1.1 of~~ [IEEE 1588], a PTP clock cancelling several types of PTP messages may use a single *Signalling* message containing multiple TLVs or multiple *Signalling* messages. Grant-ports and request-ports compliant with this profile must be prepared to handle those two situations.

The PTP clock cancelling the session may either cancel the multiple service types with a single *Signalling* message containing multiple CANCEL_UNICAST_TRANSMISSION TLVs, or with multiple independent *Signalling* messages (e.g., each containing a single CANCEL_UNICAST_TRANSMISSION TLV). The other PTP clock receiving the cancellation is allowed to respond to these requests either with a single *Signalling* message containing multiple ACKNOWLEDGE_CANCEL_UNICAST_TRANSMISSION TLVs, or with multiple independent *Signalling* messages (e.g., each containing a single ACKNOWLEDGE_CANCEL_UNICAST_TRANSMISSION TLV).

NOTE 2 – The "~~R~~renewal ~~I~~invited" flag described in [IEEE 1588] ~~clause 16.1.4.2.6~~ is not used in this profile.

6.7 Alternate BMCA, telecom slave model and master selection process

This clause describes the Alternate BMCA algorithm, the telecom slave model and the associated master selection process. These are described in the following clauses.

6.7.1 Alternate BMCA

The PTP profile specified in this Recommendation uses an Alternate BMCA, as described in clause 9.3.1 of [IEEE 1588]. This Alternate BMCA differs from the default BMCA of [IEEE 1588] in the following:

a) The Alternate BMCA considers the per-port Boolean attribute masterOnly. If masterOnly is TRUE, the port is never placed in the SLAVE state, and will always go to the MASTER state. If masterOnly is FALSE, the port can be placed in the SLAVE state. The masterOnly attribute is set via the configurable port data set member portDS.masterOnly.

The default value and range of values for this attribute, for the ports of a BC or OC that can only be a GM (i.e., T-GM), are TRUE and {TRUE}.

The default value and range of values for this attribute, for the port of a slave-only OC (i.e., T-TSC-P or T-TSC-A) are FALSE and {FALSE}.

The default value and range of values for this attribute, for the ports of a BC that may or may not be a GM (i.e., T-BC-P or T-BC-A) are TRUE and {TRUE, FALSE}.

b) The computation of E_{rbest} is according to the description provided in clause 9.3.2.3 of [IEEE 1588], with the exception that the E_{rbest} of a port r must be set to the empty set when the masterOnly attribute of this port r is set to TRUE, irrespective of any other consideration. This is so that the computation of E_{best} will not use the information contained in any *Announce* messages received on a port r where the masterOnly attribute is set to TRUE.

c) The Alternate BMCA allows for multiple clocks to be active GMs simultaneously (clocks with clockClass less than 128 cannot be a slave). If there are multiple active GMs, every clock that is not a GM is synchronized by a single GM in the PTP domain.

d) The per-port attribute localPriority is assigned to each port r of a clock and is used in the determination of E_{rbest} and E_{best}. Each parent clock or foreign master clock data set, whose *Announce* information was received on the port r, is appended with the localPriority attribute of the local port r before the data set comparison defined in Figure 3 and Figure 4 below is invoked. The localPriority attribute is not transmitted in *Announce* messages. This attribute is used as a tie-breaker in the data set comparison algorithm, in the event that all other previous attributes of the data sets being compared are equal. The localPriority attribute is set via the configurable, unsigned integer, port data set member portDS.localPriority. The data type for this attribute is UInteger8. The range of values for this attribute is {1-255}. The default value for this attribute is 128. A clock compliant with this PTP profile is allowed to support a subset of the values defined in the range.

e) The attribute localPriority is assigned to the local clock, to be used if needed when the data associated with the local clock, D_0, is compared with data on another potential GM received via an *Announce* message. The local clock localPriority attribute is set via the configurable, unsigned integer, default data set member defaultDS.localPriority. The data type for this attribute is UInteger8. The range of values for this attribute is {1-255}. The default value for this attribute is 128. A clock compliant with this PTP profile is allowed to support a subset of the values defined in the range.

f) The data set comparison algorithm is modified according to Figures 3 and 4 in clause 6.7.9.

NOTE 1 − Because the value of the masterOnly attribute is, per definition, always TRUE on all PTP ports of a T-GM, the localPriority attribute is, in practice, not used for a T-GM.

NOTE 2 − For a T-GM, the Alternate BMCA output is in practice static and provides a recommended state = BMC_MASTER, because the masterOnly attribute = TRUE for all the PTP ports of a T-GM. The resulting decision code can be M1 or M2 (see Figure 2 below), depending on the status of the T-GM (i.e., clockClass value of the T-GM).

NOTE 3 − When the value of the masterOnly attribute is TRUE on a PTP port, the PTP port typically does not request unicast services from other ports.

6.7.2 Considerations on the use of the localPriority attributes

The localPriority attributes provide a powerful tool in defining the synchronization network architecture.

The use of the default values for these attributes as defined by the Alternate BMCA results in a timing-loop free synchronization network.

Proper planning will be mandatory to avoid timing-loops when configuring values different from the default ones.

6.7.3   Static clock attribute priority1

In this PTP profile, the clock attribute priority1 is static. It is initialized to a default value equal to the midpoint value, 128, of its range, and this value must not be changed.

The priority1 parameter is not used in this version of the PTP telecom profile. Future versions may consider using this attribute, this is for further study.

6.7.4   Clock attribute priority2

In this PTP profile, the clock attribute priority2 is configurable.

It is initialized to a default value, equal for T-GM, T-BC-P, and T-BC-A clocks to the midpoint value, 128, of its range {0-255}. The default value for T-TSC-P and T-TSC-A clocks is 255, and the range is {255}.

A T-GM, T-BC-P, or T-BC-A compliant with this PTP profile must support all the values of priority2 defined in the range. A T-TSC-P or T-TSC-A compliant with this profile must support, on reception, all the values of priority2 defined in the full [IEEE 1588] range (i.e., {0-255}).

Appendix I describes possible use cases for the priority2 attribute; other cases are for further study.

6.7.5   Clock attribute clockClass

A PTP clock compliant with this PTP profile must support all values of clockClass upon reception (shall not discard) defined in the full [IEEE 1588] range. The applicable values of the clock attribute clockClass are specified in clause 6.8 of this Recommendation.

NOTE 1 – The behaviour on reception of a clockClass value not specified in Table 2 is for further study. NOTE – It is not expected that equipment compliant to this profile and deployed in an ITU-T G.8275.2 network will receive clockClass values not specified in Table 2. If a clockClass value not specified in Table 2 is received, then the equipment may raise an implementation specific alarm. Future revisions of this profile may include clockClass values not defined in Table 2.

6.7.6   Clock attribute clockAccuracy

A PTP clock compliant with this PTP profile must support all the values of clockAccuracy upon reception (shall not discard) defined in the full [IEEE 1588] range. The values that can be transmitted in the clockAccuracy field are shown in Table A.1. The following values of the clock attribute clockAccuracy apply for the following situations:

–   0x20 for a T-GM connected to an ePRTC in locked-mode (i.e., e PRTC traceable to GNSS).

–   0x21 for a T-GM connected to a PRTC in locked-mode (i.e., PRTC traceable to GNSS) or T-GM connected to ePRTC in phase/time holdover within [ITU-T G.8272.1] ePRTC-A specification as specified in Table 3 of [ITU-T G.8272.1].

–   0xFE for a T-GM not connected to an ePRTC in locked-mode, nor to an ePRTC in phase/time holdover within the [ITU-T G.8272.1] Table 3 specification, nor to a PRTC in locked mode.

–   0xFE for a T-BC-P or T-BC-A not connected to a local time reference that is traceable to GNSS in locked mode on a virtual PTP port.

The clockAccuracy for a T-BC-P or T-BC-A when connected to a local time reference that is traceable to GNSS in locked mode on a virtual PTP port is for further study.

6.7.7 Clock attribute offsetScaledLogVariance

The following values of the clock attribute offsetScaledLogVariance apply for the following situations:

– 0x4B32 for a T-GM connected to an ePRTC in locked-mode (i.e., ePRTC traceable to GNSS). This corresponds to TDEV of 10 ns, at observation interval of 1 000 000 seconds. The corresponding value of PTP Variance (PTPVAR) is 1.271×10^{-16} s^2 (see Appendix IX of [ITU-T G.8275.1]).

– 0x4E5D for a T-GM connected to a PRTC in locked-mode (i.e., PRTC traceable to GNSS) This corresponds to TDEV of 30 ns, at observation interval of 10000 seconds. The corresponding value of PTPVAR is 1.144×10^{-15} s^2 (see Appendix IX of [ITU-T G.8275.1]).

– 0xFFFF for a T-GM not connected to a PRTC in locked-mode.

– 0xFFFF for a T-BC-P or T-BC-A not connected to a local time reference that is traceable to GNSS in locked mode on a virtual PTP port.

The offsetScaledLogVariance for a T-BC-P or T-BC-A when connected to a local time reference traceable to GNSS in locked mode on a virtual PTP port is for further study.

6.7.8 State decision algorithm

The state decision algorithm applicable to the Alternate BMCA of the PTP profile specified in this Recommendation is given in Figure 2. After a decision is reached by use of this algorithm, the data sets of the local clock are updated as specified in clause 9.3.5 of [IEEE 1588]. Details on the use of the algorithm are given in clause 9.3.3 of [IEEE 1588].

6.7.9 Data set comparison algorithm

The data set comparison algorithm for the Alternate BMCA of the PTP profile specified in this Recommendation is given in Figures 3 and 4 below. With this algorithm, one clock is compared with another using the data sets representing those clocks, appended with the localPriority attribute. Details on the use of the algorithm are given in clause 9.3.4 of [IEEE 1588].

If either of the data sets, A or B, in Figures 3 and 4 contain the data of the parent clock or a foreign master clock, the corresponding localPriority for its data set is the localPriority of the local port r on which the information from that parent clock or foreign master clock has been received (see item (d) of clause 6.7.1).

If either of the data sets, A or B, in Figures 3 and 4 contain the data of the local clock, D_0, the corresponding localPriority for that data set is the localPriority of the local clock (see item (e) of clause 6.7.1).

NOTE 1 – It is recommended that the entire data set comparison algorithm described in Figures 3 and 4 be implemented even if some parameters are currently static, because they may be used in future versions of this Recommendation.

NOTE 2 – If portDS.SF is TRUE on port r, then the PTP port should set the respective E_{rbest} to the empty set. As a result, the computation of E_{best} will not use the information contained in any Announce messages received on the port r. Signal fail (SF) is described in clause 6.7.11.
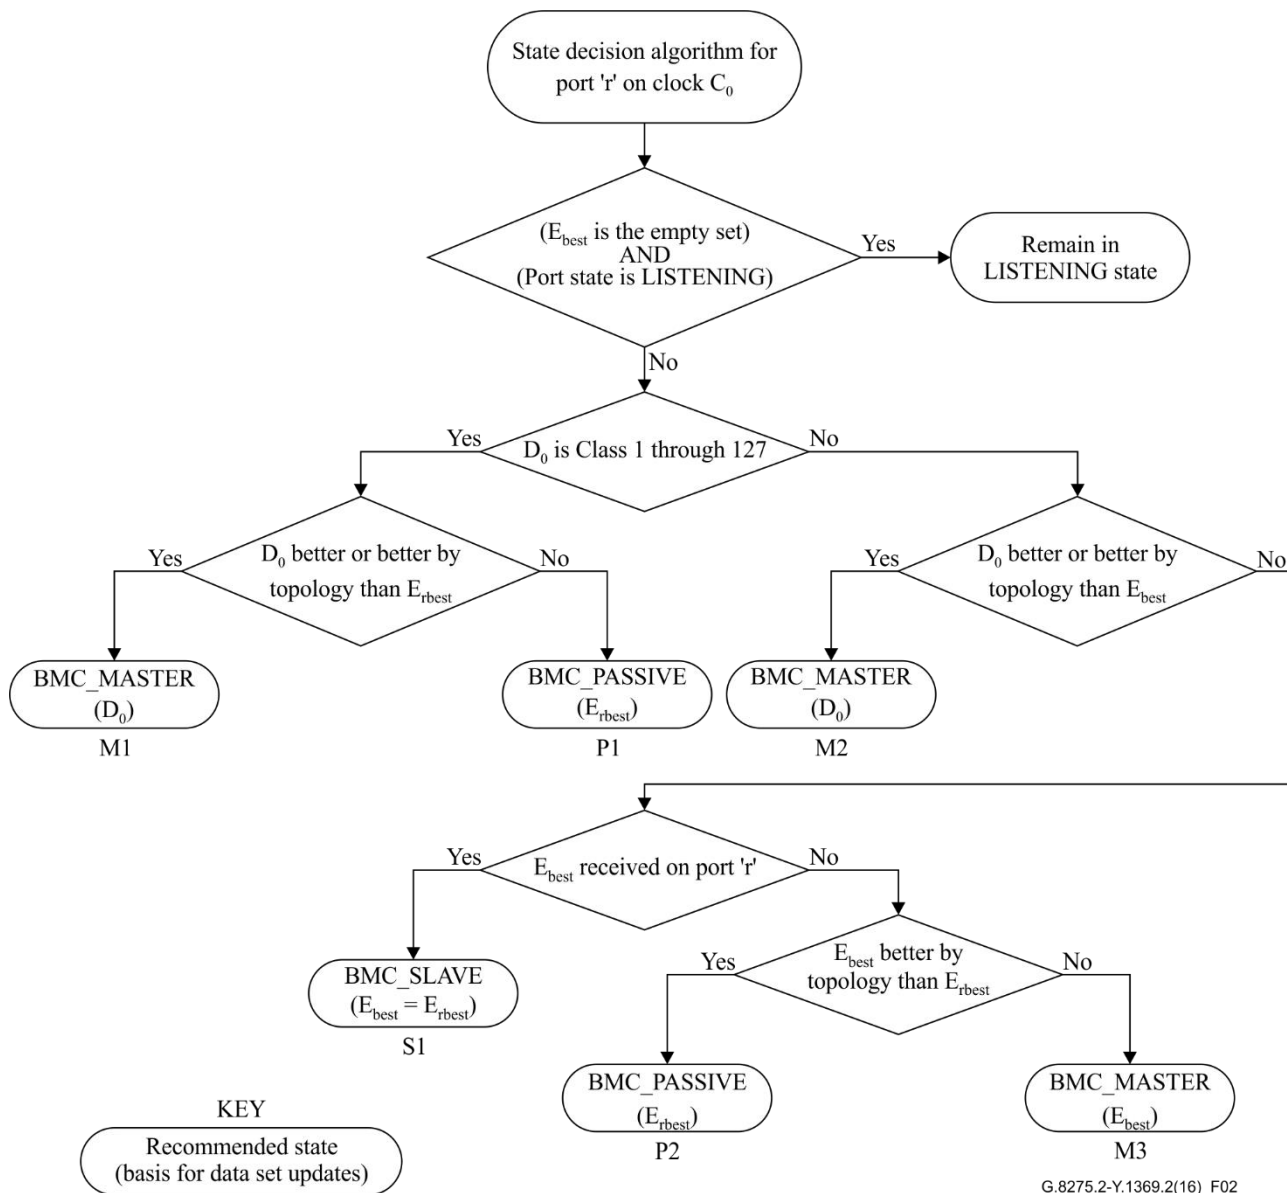
Figure 2 – State decision algorithm for Alternate BMCA
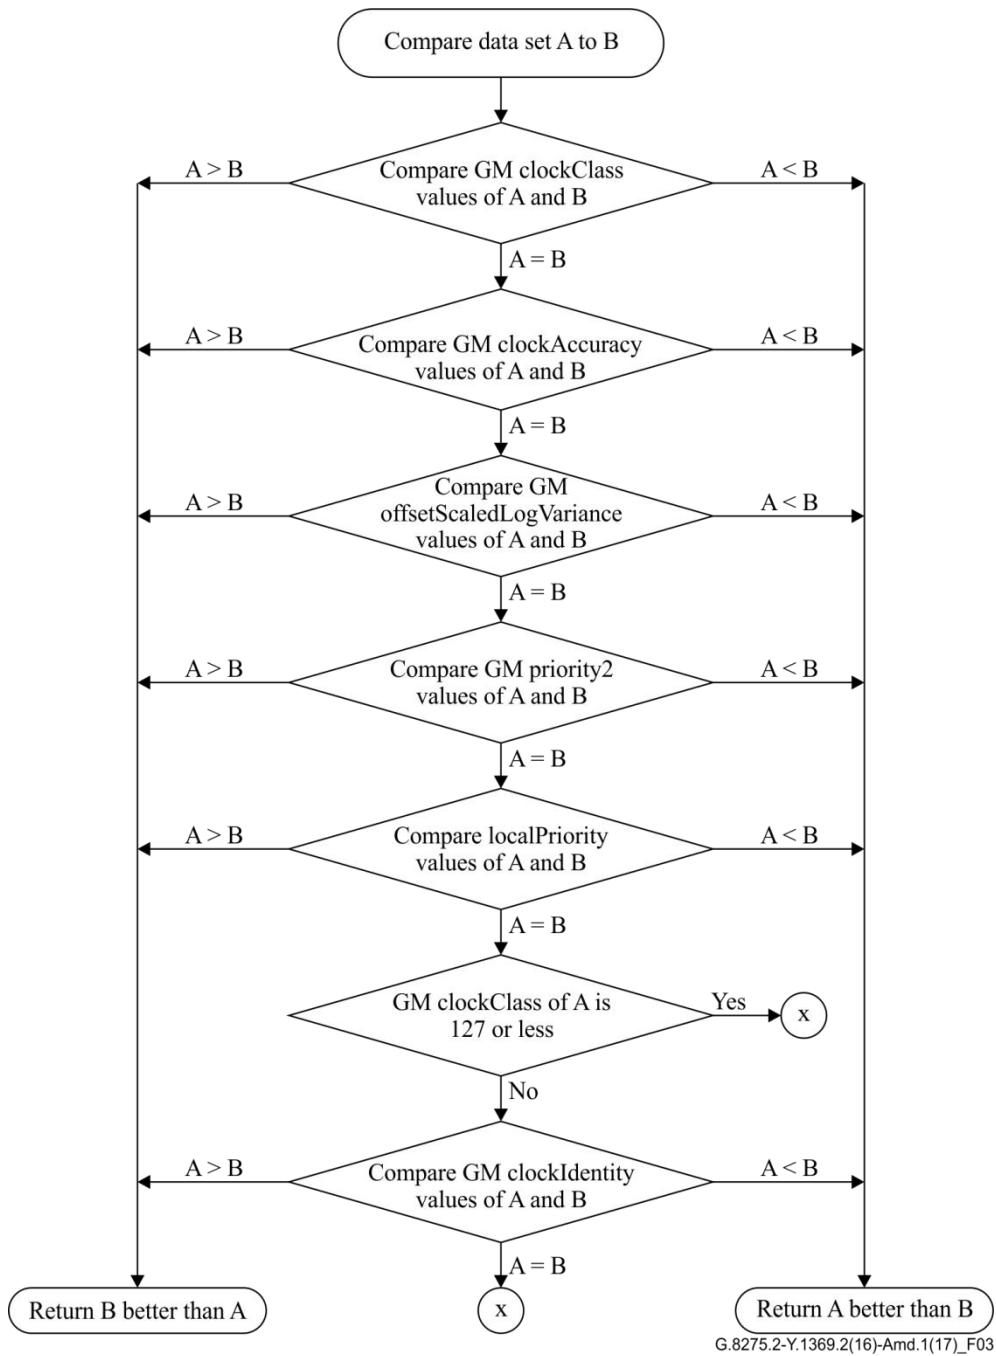
Figure 3 – Data set comparison algorithm, part 1, for Alternate BMCA
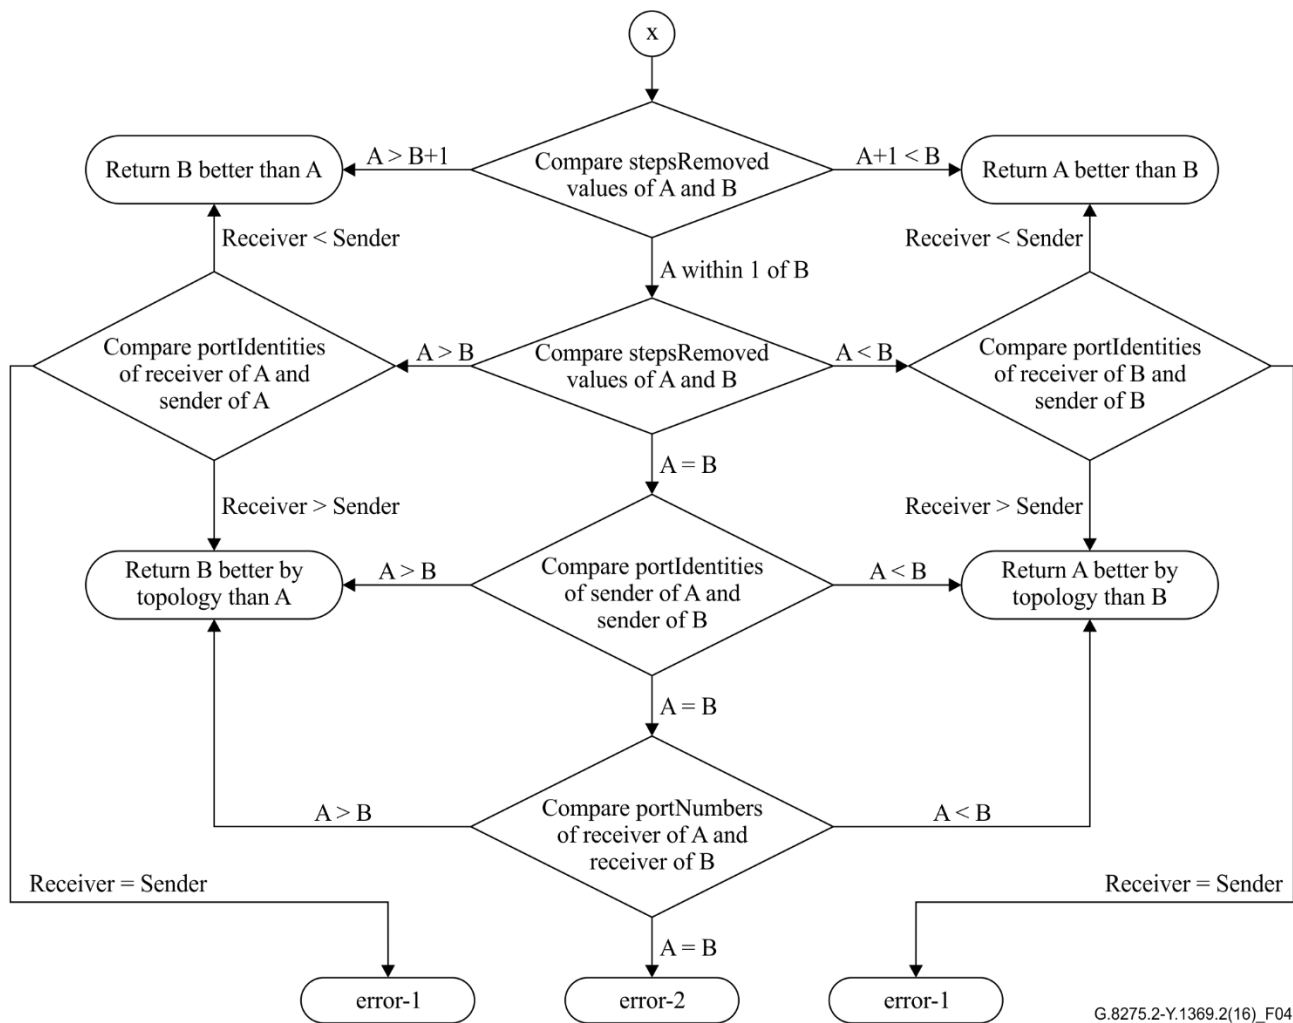
Figure 4 – Data set comparison algorithm, part 2, for Alternate BMCA

NOTE 3 – stepsRemoved used in the BMCA does not characterize or reflect the amount of packet delay variation (PDV) or asymmetry on a connection. The BMCA may not select the path with the lowest PDV or asymmetry.

6.7.10 Unused PTP fields

Some PTP fields are not used in this PTP profile. This clause defines the actions applicable to these unused PTP fields.

Table A.6 in clause A.10 defines the PTP common header flag values, and whether or not each flag is used in this profile.

In addition, the following fields are not used in this profile:

– The "controlField" in the common header of PTP messages is not used in this profile. This field must be ignored by the receiver for all types of PTP messages.

– The "priority1" field in the *Announce* message is not used and must be set to a fixed value specified in clause 6.7.3.

When a PTP clock receives a PTP message with a field, whose use is not specified in this PTP profile, containing a value outside the allowed range, then this field of the PTP message must be ignored, without discarding the PTP message.

As an example, a PTP clock compliant with this PTP profile must ignore on reception the field value for the following fields. A clock compliant with this PTP profile must not update its local data sets with the ingress value for these fields.

–   flagField – PTP profile Specific 1

–   flagField – PTP profile Specific 2

When a PTP clock receives a PTP message with a field, whose use is specified in this PTP profile, containing a value outside the allowed range for reception, then this entire PTP message must be discarded. The allowed range for a parameter on reception is the same as the range for the corresponding default dataset parameter, except for the attributes clockClass, clockAccuracy, offsetScaledLogVariance, and priority2 (see clauses 6.7.4, 6.7.5, 6.7.6 and 6.7.7).

As an example, a compliant clock must discard on reception the ingress packet (General and Event messages) when any of the following fields are outside of the allowed range for the profile. The clock's local data set must not be updated with the ingress value.

–   domainNumber

–   versionPTP

–   flagField – unicastFlag

NOTE 1 − If a clock receives an *Announce* message with the "priority1" field set to a value other than 128, and if the clock advertising this value is selected as the GM, then 128 must be re-advertised by the receiving clock. The unused attribute priority1 is ignored by the receiving clock for the purpose of the Alternate BMCA.

NOTE 2 − The allowed ranges for reception for the clock attributes priority2, clockClass, clockAccuracy, and offsetScaledLogVariance are the respective full [IEEE 1588] ranges, see clauses 6.7.4, 6.7.5, 6.7.6 and 6.7.7 of this Recommendation.

6.7.11   Packet timing signal fail

This clause is optional but, if implemented, it is necessary for the equipment to conform to the requirements contained herein. An implementation may support all, none, or a subset of these types of PTSF.

This clause defines the notion of packet timing signal fail (PTSF), which corresponds to a signal indicating a failure of the PTP packet timing signal received by the slave.

Three types of PTSF may be raised in a PTP clock:

1)   PTSF-lossSync, lack of reception of PTP timing messages from a grant-port(loss of the packet timing signal): if the request-port no longer receives the timing messages sent by a grant-port (i.e., *Sync* and subsequently *Follow_up Follow_Up* and *Delay_Resp* messages), then a PTSF-lossSync associated to with this grant-port must occur. A timeout period for reception of *Sync* messages or *Delay_Resp* messages (i.e., syncReceiptTimeout and delayRespReceiptTimeout) for these timing messages must be implemented in the request-port before triggering the PTSF-lossSync (the range and default value of this these timeout parameters are defined in Table A.5).

The value of syncReceiptTimeout shall specify the number of *Sync* message intervals that have to pass without receipt of a *Sync* and, if the twoStep flag of the Sync message is TRUE, or a *Follow_Up* message before the triggering of the PTSF-lossSync event.

The value of delayRespReceiptTimeout shall specify the number of *Delay_Req* message intervals that have to pass without receipt of a *Delay_Resp* message before the triggering of the PTSF-lossSync event.

A timer should be activated after receiving the grant message from the grant port. If the request port sends a new signalling message to change message rate, it shall terminate the timer. Once the request port receives the new grant message, the associated timer shall restart.

2)   PTSF-unusable, unusable PTP packet timing signal received by the request-port, exceeding the input tolerance of the request-port (noisy packet timing signal): if the PTP packet timing signal is not usable for the request-port to achieve the performance target (e.g., violates the request-port input tolerance because of excessive PDV noise), then a PTSF-unusable associated ~~to~~ with this master must occur. The criteria used to determine that the packet timing signal is not suitable to be used ~~is~~ are for further study (an example of criteria to be studied may relate to the PDV experienced by the packet timing signal as it traverses the network from the grant-port to the request-port).

3)   PTSF-syncUncertain, uncertain timing signal received by the request-port: if the synchronizationUncertain flag of an *Announce* message received from a grant-port is TRUE, a PTSF-syncUncertain associated ~~to~~ with this grant-port must occur.

When a PTSF occurs, the clock ~~will~~ may set the PTP portDS.SF to TRUE and generate a state decision event. As described in clause 6.7.9 a value of portDS.SF of TRUE can be used to exclude PTP ports from the BMCA selection process. An implementation may set the PTP portDS.SF to TRUE for only a subset of PTSF listed above.

NOTE – See Appendix V and Appendix VII for further information.

6.8   Phase/time traceability information

In order to deliver phase/time traceability information, the clockClass values described in Table 2 below must be used in this PTP telecom profile.

The frequencyTraceable flag present in the header of the PTP messages is defined in this profile as follows: if the PTP clock is traceable to a PRTC in locked mode or to a primary reference clock (PRC), e.g., using a PRC-traceable physical layer frequency input, then this parameter must be set to TRUE, otherwise it must be FALSE. This flag is not used in the Alternate BMCA defined in clause 6.7; the values provided for this flag in Table 2 can be used by the network operator for monitoring purposes or by the end applications to take definitive action as described in Appendix II.

When a T-GM first enters holdover, it downgrades the clockClass value that it uses to 7. It then calculates if the time error at its output is still within the holdover specification. When the T-GM determines that the time error at its output has exceeded the holdover specification, it downgrades the clockClass value that it uses to 140, 150 or 160 depending on the quality of its frequency reference (internal oscillator or physical layer frequency signal received on an external interface).

As an example, when a T-BC-P or T-BC-A first enters holdover, it downgrades the clockClass value that it uses to 135. It then calculates if the time error at its output is still within the holdover specification. When the T-BC-P or T-BC-A determines that the time error at its output has exceeded the holdover specification, it downgrades the clockClass value that it uses to 165 (internal oscillator or received physical layer frequency signal on an external interface).

NOTE 1 – The applicable holdover specification depends on the design and budgeting of the synchronization network.

NOTE 2 – The case of a T-BC-P or T-BC-A acting as a GM, with an external phase/time input coming from a PRTC, is handled by means of a virtual PTP port with associated E_{rbest} attributes as described in Annex C of this Recommendation. The general case of a T-BC-P or T-BC-A with a phase/time external synchronization input different from PRTC is for further study.

NOTE 3 – For the T-BC-P or T-BC-A in locked mode, the traceability information of the currently selected best master clock will be passed to the downstream nodes, as per PTP. This means that the attributes and flags in the PTP header will always reflect the phase/time traceability information from the current parent clock, regardless of the frequency traceability of the T-BC-P/A's physical layer clock. Failure scenarios including holdover are for further study.

Table 2 – Applicable clockClass values

Phase/time traceability description	defaultDS clockClass	frequencyTraceable flag	timeTraceable flag
T-GM connected to a PRTC in locked mode (e.g., PRTC traceable to GNSS)	6	TRUE	TRUE
T-GM in holdover, within holdover specification, traceable to Category 1 frequency source (Note 1)	7	TRUE	TRUE
T-GM in holdover, within holdover specification, non-traceable to Category 1 frequency source (Note 1)	7	FALSE	TRUE
T-BC-P or T-BC-A in holdover, within holdover specification, traceable to Category 1 frequency source (Note 1)	135	TRUE	TRUE
T-BC-P or T-BC-A in holdover, within holdover specification, non-traceable to Category 1 frequency source (Note 1)	135	FALSE	TRUE
T-GM in holdover, out of holdover specification, traceable to Category 1 frequency source (Note 1)	140	TRUE	FALSE
T-GM in holdover, out of holdover specification, traceable to Category 2 frequency source (Note 1)	150	FALSE	FALSE
T-GM in holdover, out of holdover specification, traceable to Category 3 frequency source (Note 1)	160	FALSE	FALSE
T-BC-P or T-BC-A in holdover, out of holdover specification (Note 1)	165	(Note 2)	FALSE
T-GM, T-BC-P, T-BC-A, in free-run mode, or, T-TSC-P or T-TSC-A, acting as a BC in free-run mode	248	(Note 2)	FALSE
T-GM, T-BC-P, T-BC-A, without time reference since start-up, or, T-TSC-P or T-TSC-A, acting as a BC without time reference	248	(Note 2)	FALSE
T-TSC-P or T-TSC-A acting as an OC	255	(Note 2)	As per PTP
NOTE 1 – The holdover specification threshold controlling the time spent advertising clockClass values 7 or 135 could be set to zero so that the T-GM, T-BC-P, or T-BC-A would advertise a degraded clockClass value directly after losing traceability to a PRTC. In this case, initially after advertising clockClass values			

Table 2 – Applicable clockClass values

Phase/time traceability description	defaultDS clockClass	frequencyTraceable flag	timeTraceable flag
140, 150, 160 or 165, a clock may still be within the holdover specification. For a description of frequency source "Category" see Table 3 below. NOTE 2 – The frequencyTraceable flag may be TRUE or FALSE, depending on the availability of a PRC-traceable physical layer frequency input signal. NOTE 3 – The term "holdover" in this table refers to "time holdover".			

Table 3 describes how the clock quality levels (QLs) defined in [ITU-T G.781] are mapped to Category 1, 2, and 3 frequency sources used in Table 2.

Table 3 – Mapping of ITU-T G.781 clock QLs to Category 1, 2, 3 frequency sources

Category (in Table 2)	ITU-T G.781 Option I QLs	ITU-T G.781 Option II QLs
Category 1 frequency source	QL-PRC	QL-PRS
Category 2 frequency source	QL-SSU-A	QL-ST2
Category 3 frequency source	QL-SSU-B	QL-ST3E
NOTE – Other frequency source categories, while not used in Table 2, are possible. An example is a category containing QL-EEC1 and QL-EEC2.		

6.9 Use of alternate master flag

A PTP clock must only synchronize to a PTP timing service being provided by its parent clock, whose port is in the PTP MASTER state. To ensure this operation, this profile uses the alternateMasterFlag field defined in clause 7.3.8.2 of [IEEE 1588] with the following behaviour.

a) On transmission of an Announce message, a PTP port will set the alternateMasterFlag to 0 when the transmitting PTP port state is MASTER; otherwise the PTP port must set the alternateMasterFlag to 1.

b) Referring to clause 13.3.2.6 of [IEEE 1588-2008] and clause 13.3.2.8 of [IEEE 1588-2019], the alternateMasterFlag is only set on transmission of Announce, Sync, Follow_Up and Delay_Resp messages.

c) On reception, a PTP port that receives a PTP Announce message with alternateMasterFlag value 1 must discard (and not process) the message. For example, such an Announce message must not be input into the BMCA.

While the alternateMasterFlag is used in this version of the profile, clause 17.4 of [IEEE 1588-2008] and clause 17.3 of [IEEE 1588-2019] is are not used.

7 ITU-T PTP profile for phase/time distribution with partial timing support from the network

The [IEEE 1588] profile that supports time distribution in unicast mode is contained in Annex A.

8 Security aspects

Security aspects are for further study.

Annex A

ITU-T PTP profile for time distribution with partial timing support
from the network (unicast mode)

(This annex forms an integral part of this Recommendation.)

This annex contains the telecom profile for time distribution as required by [IEEE 1588]. In order to claim compliance with the telecom profile, the requirements in this annex and in the body of this Recommendation must both be met.

A.1   Profile identification

profileName: ITU-T PTP profile for time distribution with partial timing support from the network (unicast mode)

profileVersion: 1.1 (for an implementation based on IEEE Std 1588-2008 [IEEE 1588-2008])

profileVersion: 1.2 (for an implementation based on IEEE Std 1588-2019 [IEEE 1588-2019])

profileIdentifier: 00-19-A7-02-01-01 (for an implementation based on IEEE Std 1588-2008 [IEEE 1588-2008])

profileIdentifier: 00-19-A7-02-01-02 (for an implementation based on IEEE Std 1588-2019 [IEEE 1588-2019])

NOTE – Version 1.1 is backward compatible with version 1.0. Equipment with version 1.0 may be deployed in the same network as equipment with version 1.0. Version 1.1 adds additional optional functionality that is not present in version 1.0. Version 1.2 is backward compatible with version 1.1 and version 1.0. Equipment with version 1.2 may be deployed in the same network as equipment with version 1.1 or version 1.0. Version 1.2 indicates an implementation compliant with IEEE1588-2019. Versions 1.1 and 1.2 are both valid versions for this edition of this Recommendation, i.e., Version 1.2, based on [IEEE 1588-2019] does not supersede Version 1.1, based on [IEEE 1588-2008].

This profile is specified by ITU-T.

A copy may be obtained from www.itu.int.

A.2   PTP attribute values

The default values and ranges of the PTP attributes for use in this profile are contained in Tables A.1, A.2, A.3, A.4 and A.5. For the attributes clockClass, clockAccuracy, offsetScaledLogVariance, and priority2, the ranges shown are those for the defaultDS.

NOTE – A boundary clock follows the rules of [IEEE 1588] for selection of parent clock, updating of parentDS, and transmission of Announce messages, so it may transmit values different from the defaultDS values.

Attributes not specified by this profile shall use the [IEEE 1588] default initialization values and ranges.

Some   attributes   in   these   tables   are   associated   with   optional   features   of   this Recommendation. Therefore, these attributes are also optional unless the associated feature is implemented, in which case these attributes must be supported.

Table A.1 – defaultDS data set member specifications

Clause from [IEEE 1588-2008]	Clause from [IEEE 1588-2019]	Members of the data set	T-GM requirements		T-TSC-P and T-TSC-A requirements		T-BC-P and T-BC-A requirements	
			Default initialization value	Range	Default initialization value	Range	Default initialization value	Range
8.2.1.2.1	(Note 5)	defaultDS.twoStepFlag (static)	As per PTP	{FALSE, TRUE}	As per PTP	{FALSE, TRUE}	As per PTP	{FALSE, TRUE}
8.2.1.2.2	8.2.1.2.2	defaultDS.clockIdentity (static)	As per PTP, based on EUI-64 format	As per PTP	As per PTP, based on EUI-64 format	As per PTP	As per PTP, based on EUI-64 format	As per PTP
8.2.1.2.3	8.2.1.2.3	defaultDS.numberPorts (static)	1 for OC As per PTP for BC	{1} for OC As per PTP for BC	1 for OC As per PTP for BC	{1} for OC As per PTP for BC	As per PTP	As per PTP
8.2.1.3.1.1	8.2.1.3.1.2	defaultDS.clockQuality.clockClass (dynamic)	248	{6, 7, 140, 150, 160, 248}	255 for OC 248 for BC	{255} for OC {248} for BC	248	{135, 165, 248}
8.2.1.3.1.2	8.2.1.3.1.3	defaultDS.clockQuality.clockAccuracy (dynamic)	0xFE (Note 2)	As per PTP (Note 2) (Note 4)	0xFE (Note 2)	{0xFE} (Note 2)	0xFE (Note 2)	{0xFE} (Note 2)
8.2.1.3.1.3	8.2.1.3.1.4	defaultDS.clockQuality.offsetScaledLogVariance (dynamic)	0xFFFF	As per PTP (Note 4)	0xFFFF	{0xFFFF}	0xFFFF	{0xFFFF}
8.2.1.4.1	8.2.1.4.1	defaultDS.priority1 (configurable)	128 (Note 1)	{128} (Note 1)	128 (Note 1)	{128} (Note 1)	128 (Note 1)	{128} (Note 1)
8.2.1.4.2	8.2.1.4.2	defaultDS.priority2 (configurable)	128	{0-255}	255	{255}	128	{0-255}
8.2.1.4.3	8.2.1.4.3	defaultDS.domainNumber (configurable)	44	{44-63}	44	{44-63}	44	{44-63}
8.2.1.4.4	8.2.1.4.4	defaultDS.slaveOnly (configurable)	FALSE	{FALSE}	TRUE for OC FALSE for BC	{TRUE} for OC {FALSE} for BC	FALSE	{FALSE}
(Note 6)	8.2.1.4.5	defaultDS.sdoId (configurable)	0x000	0x000	0x000	0x000	0x000	0x000
New member	New member	defaultDS.localPriority (configurable)	128	{1-255}	128	{1-255}	128	{1-255}
New member	New member	defaultDS.SF (dynamic)	FALSE	{FALSE}	FALSE	{FALSE}	FALSE	{FALSE}

NOTE 1 – As per PTP, not applicable for this profile.

NOTE 2 – For the case where the PTP grandmaster is syntonized to a PRC for frequency, but not synchronized to a reference source of time, the grandmaster should set defaultDS.clockQuality.clockAccuracy to 0xFE, "UNKNOWN".

NOTE 3 – Equipment implementing multiple slave ports, with defaultDS.clockClass value of 255, should be treated as having multiple instantiations of slave-only OCs. This is out of scope of this Recommendation.

NOTE 4 – Examples of applicable values are shown in clauses 6.7.6 and 6.7.7.

NOTE 5 – In the case of [IEEE 1588-2019] based implementations, this data set member is deprecated.

NOTE 6 – Applicable to [IEEE 1588-2019] based implementations; this data set member does not exist in [IEEE 1588-2008]

Table A.2 – currentDS data set member specifications

Clause from [IEEE 1588-2008]	Clause from [IEEE 1588-2019]	Members of the data set	T-GM requirements		T-TSC-P and T-TSC-A requirements		T-BC-P and T-BC-A requirements	
			Default initialization value	Range	Default initialization value	Range	Default initialization value	Range
8.2.2.2	8.2.2.2	currentDS.stepsRemoved (dynamic)	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP
8.2.2.3	8.2.2.3	currentDS.offsetFromMaster (dynamic)	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP
8.2.2.4	(Note 1)	currentDS.meanPathDelay (dynamic)	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP
(Note 2)	8.2.2.4	currentDS.meanDelay (dynamic)	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP
NOTE 1 – In the case of [IEEE 1588-2019] based implementations this data set member is deprecated. NOTE 2 – Applicable to [IEEE 1588-2019] based implementations; this data set member does not exist in [IEEE 1588-2008].								

Table A.3 – parentDS data set member specifications

Clause from [IEEE 1588-2008]	Clause from [IEEE 1588-2019]	Members of the data set	T-GM requirements		T-TSC-P and T-TSC-A requirements		T-BC-P and T-BC-A requirements	
			Default initialization value	Range	Default initialization value	Range	Default initialization value	Range
8.2.3.2	8.2.3.2	parentDS.parentPortIdentity (dynamic)	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP
8.2.3.3	8.2.3.3	parentDS.parentStats (dynamic)	(Note)	(Note)	(Note)	(Note)	(Note)	(Note)
8.2.3.4	8.2.3.4	parentDS.observedParentOffsetScaledLogVariance (dynamic)	(Note)	(Note)	(Note)	(Note)	(Note)	(Note)
8.2.3.5	8.2.3.5	parentDS.observedParentClockPhaseChangeRate (dynamic)	(Note)	(Note)	(Note)	(Note)	(Note)	(Note)
8.2.3.6	8.2.3.6	parentDS.grandmasterIdentity (dynamic)	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP
8.2.3.7	8.2.3.7	parentDS.grandmasterClockQuality (dynamic)	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP
8.2.3.8	8.2.3.8	parentDS.grandmasterPriority1 (dynamic)	As per PTP (Note)	As per PTP (Note)	As per PTP (Note)	As per PTP (Note)	As per PTP (Note)	As per PTP (Note)
8.2.3.9	8.2.3.9	parentDS.grandmasterPriority2 (dynamic)	As per PTP (Note)	As per PTP (Note)	As per PTP (Note)	As per PTP (Note)	As per PTP (Note)	As per PTP (Note)
NOTE – As per PTP, not applicable for this profile.								

Table A.4 – timePropertiesDS data set member specifications

Clause from [IEEE 1588-2008]	Clause from [IEEE 1588-2019]	Members of the data set	T-GM requirements		T-TSC-P and T-TSC-A requirements		T-BC-P and T-BC-A requirements	
			Default initialization value	Range	Default initialization value	Range	Default initialization value	Range
8.2.4.2	8.2.4.2	timePropertiesDS.currentUtcOffset (dynamic)	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP
8.2.4.3	8.2.4.3	timePropertiesDS.currentUtcOffsetValid (dynamic)	FALSE	{FALSE, TRUE}	FALSE	{FALSE, TRUE}	FALSE	{FALSE, TRUE}
8.2.4.4	8.2.4.4	timePropertiesDS.leap59 (dynamic)	FALSE	{FALSE, TRUE}	FALSE	{FALSE, TRUE}	FALSE	{FALSE, TRUE}
8.2.4.5	8.2.4.5	timePropertiesDS.leap61 (dynamic)	FALSE	{FALSE, TRUE}	FALSE	{FALSE, TRUE}	FALSE	{FALSE, TRUE}
8.2.4.6	8.2.4.6	timePropertiesDS.timeTraceable (dynamic)	FALSE	{FALSE, TRUE}	FALSE	{FALSE, TRUE}	FALSE	{FALSE, TRUE}
8.2.4.7	8.2.4.7	timePropertiesDS.frequencyTraceable (dynamic)	FALSE	{FALSE, TRUE} (Note)	FALSE	{FALSE, TRUE} (Note)	FALSE	{FALSE, TRUE} (Note)
8.2.4.8	8.2.4.8	timePropertiesDS.ptpTimescale (dynamic)	TRUE	{TRUE}	TRUE	{TRUE}	TRUE	{TRUE}
8.2.4.9	8.2.4.9	timePropertiesDS.timeSource (dynamic)	0xA0	As per PTP	0xA0	As per PTP	0xA0	As per PTP

NOTE – If the clock is traceable to a PRTC in locked mode or a PRC (e.g., using a PRC-traceable physical layer frequency input), then this parameter must be set to TRUE, otherwise it must be FALSE.

Table A.5 − portDS data set member specifications

Clause from [IEEE 1588-2008]		Clause from [IEEE 1588-2019]		Members of the data set	Master port requirements of T-GM		Slave port requirements of T-TSC-P and T-TSC-A		T-BC-P and T-BC-A requirements	
Clause	Data type	Clause	Data type		Default initialization value	Range	Default initialization value	Range	Default initialization value	Range
8.2.5.2.1	As per PTP	8.2.15.2.1	As per PTP	portDS.portIdentity.clockIdentity (static)	As per PTP, based on EUI-64 format	As per PTP	As per PTP, based on EUI-64 format	As per PTP	As per PTP, based on EUI-64 format	As per PTP

Table A.5 – portDS data set member specifications

Clause from [IEEE 1588-2008]		Clause from [IEEE 1588-2019]		Members of the data set	Master port requirements of T-GM		Slave port requirements of T-TSC-P and T-TSC-A		T-BC-P and T-BC-A requirements	
Clause	Data type	Clause	Data type		Default initialization value	Range	Default initialization value	Range	Default initialization value	Range
8.2.5.2.1	As per PTP	8.2.15.2.1	As per PTP	portDS.portIdentity.portNumber (static)	1 for OC As per PTP for BC	{1} for OC As per PTP for BC	1 for OC As per PTP for BC	{1} for OC As per PTP for BC	As per PTP	As per PTP
8.2.5.3.1	As per PTP	8.2.15.3.1	As per PTP	portDS.portState (dynamic)	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP	As per PTP
8.2.5.3.2	As per PTP	8.2.15.3.2	As per PTP	portDS.logMinDelayReqInterval (dynamic)	(Note 1)	(Note 1)	(Note 1)	(Note 1)	(Note 1)	(Note 1)
8.2.5.3.3	As per PTP	(Note 5)		portDS.peerMeanPathDelay (dynamic)	(Note 1)	(Note 1)	(Note 1)	(Note 1)	(Note 1)	(Note 1)
8.2.5.4.1	As per PTP	8.2.15.4.1	As per PTP	portDS.logAnnounceInterval (configurable)	(Note 1)	(Note 1)	(Note 1)	(Note 1)	(Note 1)	(Note 1)
8.2.5.4.2	As per PTP	8.2.15.4.2	As per PTP	portDS.announceReceiptTimeout (configurable)	2	{2}	As per PTP	As per PTP	As per PTP	As per PTP
8.2.5.4.3	As per PTP	8.2.15.4.3	As per PTP	portDS.logSyncInterval (configurable)	(Note 1)	(Note 1)	(Note 1)	(Note 1)	(Note 1)	(Note 1)
8.2.5.4.4	As per PTP	8.2.15.4.4	As per PTP	portDS.delayMechanism (configurable)	01 (Note 2)	{01} (Note 2)	'01' for a two-way slave-port, and 'FE' for a one-way slave-port	{01,FE}	01	{01}
8.2.5.4.5	As per PTP	8.2.15.4.5	As per PTP	portDS.logMinPdelayReqInterval (configurable)	(Note 1)	(Note 1)	(Note 1)	(Note 1)	(Note 1)	(Note 1)
8.2.5.4.6	As per PTP	8.2.15.4.6	As per PTP	portDS.versionNumber (configurable)	2	{2}	2	{2}	2	{2}
(Note 7)		8.2.15.4.7	As per PTP	portDS.minorVersionNumber (configurable)	1	{1}	1	{1}	1	{1}
New member	Boolean	New member	Boolean	portDS.masterOnly (configurable)	TRUE	{TRUE}	FALSE	{FALSE}	TRUE	{TRUE, FALSE}
New member	UInteger8	New member	UInteger8	portDS.localPriority	128	{1-255}	128	{1-255}	128	{1-255}

Table A.5 – portDS data set member specifications

Clause from [IEEE 1588-2008]		Clause from [IEEE 1588-2019]		Members of the data set	Master port requirements of T-GM		Slave port requirements of T-TSC-P and T-TSC-A		T-BC-P and T-BC-A requirements	
Clause	Data type	Clause	Data type		Default initialization value	Range	Default initialization value	Range	Default initialization value	Range
				(configurable)						
New member	UInteger16	New member	UInteger16	portDS.SF (dynamic)	FALSE	{FALSE}	FALSE	{TRUE, FALSE}	FALSE	{TRUE, FALSE}
New Member	UInteger16	New Member	UInteger16	portDS.syncReceiptTimeout (configurable)	FFSNA	FFSNA	FFS (Note 3) (Note 4)	FFS {3 – 65535} (Note 6)	FFS (Note 3)	FFS {3 – 65535} (Note 6)
New Member	Boolean	New Member	Boolean	portDS.delayRespReceiptTimeout (configurable)	FFSNA	FFSNA	FFS (Note 3) (Note 4)	FFS {3 – 65535} (Note 6)	FFS (Note 3)	FFS {3 – 65535} (Note 6)

NOTE 1 – As per PTP, not applicable for this profile.

NOTE 2 – The master must support two-way operation.

NOTE 3 – Implementation specific.

NOTE 4 – An implementation can choose a fixed value or base it on negotiated message rate (e.g., a value proportional to the message rate). See Appendix X for further details.

NOTE 5 – In case of [IEEE 1588-2019] based implementations, this data set member is deprecated.

NOTE 6 – The full range is not expected to be supported by an implementation, as the receipt timeout value typically depends on the message rate and the ability of the PTP clock implementation to maintain frequency and time during loss of Sync and/or Delay_Resp messages. An implementation should be verified only over the expected operating conditions. See Appendix X

NOTE 7 – Applicable to [IEEE 1588-2019] based implementations; this data set member does not exist in [IEEE 1588-2008].

A.3 PTP options

A.3.1 Node types required, permitted or prohibited

In this profile, the permitted node types are ordinary clocks and boundary clocks.

The use of transparent clocks is for further study.

A.3.2 Transport mechanisms required, permitted, or prohibited

In this profile, the required transport mechanism is *Transport of PTP over User Datagram Protocol over Internet Protocol Version 4* UDP/IPv4 as per Annex D in [IEEE 1588]. Bit 0 of the transportSpecific field defined in [IEEE 1588-2008] must be set to "0"; that field does not exist in [IEEE 1588-2019].

In this profile, a permitted transport mechanism is *Transport of PTP over User Datagram Protocol over Internet Protocol Version 6* UDP/IPv6 as per Annex E in [IEEE 1588].

A.3.3 Unicast messages

All messages are sent in unicast.

In this telecom profile, unicast negotiation is enabled per default.

The slave will initiate the session by following the unicast message negotiation procedure defined in [IEEE 1588] clause 16.1.

A.3.4 REQUEST_UNICAST_TRANSMISSION TLV

The value of logInterMessagePeriod is the logarithm, to base 2, of the requested mean period, in seconds, between the requested unicast messages.

For requesting unicast *Announce* messages: The configurable range is 0 to –3 (which represents a range from 1 message per second to eight messages per second). No default rate is specified.

For requesting unicast *Sync* messages: The configurable range is 0 to –7 (which represents a range from 1 message per second to 128 messages per second). No default rate is specified.

For requesting unicast *Delay_Resp* messages: The configurable range is 0 to –7 (which represents a range from 1 message per second to 128 messages per second). No default rate is specified.

The durationField value in each REQUEST_UNICAST_TRANSMISSION TLV has a default initialization value of 300 seconds. The configurable range is 60 seconds to 1000 seconds.

NOTE 1 – A specific slave implementation, in order to meet its target performance requirements, as normal operation, may support a subset of the message rates within the ranges noted above. A master, on the other hand, is required to support the full range of message transmission rates. Unless an implementation specifies otherwise, the default value listed above is assumed to be used.

NOTE 2 – A specific slave implementation may support a subset of the durationField values within the range noted above. A master, on the other hand, is required to support the full range of durationField values. Unless an implementation specifies otherwise, the default value listed above is assumed to be used.

The maintenance and configuration of these default and configuration range values is implementation specific.

A.3.5 GRANT_UNICAST_TRANSMISSION TLV

In implementing the GRANT_UNICAST_TRANSMISSION TLV mechanism, the granted values shall be the same as requested in the received REQUEST_UNICAST_TRANSMISSION TLV as long as the requests are in the configurable range.

A.4 Best master clock algorithm options

This profile uses the Alternate BMCA described in clause 6.7 of this Recommendation.

A.5 Path delay measurement option (delay request/delay response)

The delay request/delay response mechanism can be used in this profile. The peer delay mechanism shall not be used in this profile.

A.6 Configuration management options

Management aspects are for further study and will be specified in a future version of this profile.

A.7 Clock identity format

~~For implementations compliant to [IEEE 1588-2008], the~~ The use of [IEEE EUI-64] to generate the ~~clock identity must be supported as indicated in clause 7.5.2.2.2 of [IEEE 1588]. Non-IEEE extended unique identifier (EUI) formats are not supported.~~

For implementations based on [IEEE 1588-2008], the procedures to use an EUI-48 to create the EUI-64 clockIdentity as described in clause 7.5.2.2.2 of [IEEE 1588-2008] are no longer recommended. If a clockIdentity is formed by mapping an EUI-48 to an EUI-64, and if the EUI-48 was assigned from an MA-M or MA-S, it is possible that the clockIdentity will be a duplicate of a clockIdentity formed directly from a different MA-M or MA-S (i.e., by appending bits to the end of that different MA-M or MA-S). Only if the EUI-48 was formed from an OUI (MA-L), is the uniqueness ensured. For new implementations based on [IEEE 1588-2008], the clockIdentity shall be constructed as per clause 7.5.2.2.2 of [IEEE 1588-2019]. Non-IEEE clockIdentity formats are not supported.

For implementations compliant to [IEEE 1588-2019], the clockIdentity shall be constructed as per clause 7.5.2.2.2 of [IEEE 1588-2019].

A.8   Security aspects

Security aspects are for further study. ~~The experimental security protocol of Annex K of [IEEE 1588-2008] is not used.~~

A.9   Other optional features of IEEE 1588

Other optional features of [IEEE 1588] are not used in this version of the profile.

For implementations compliant to [IEEE 1588-2008] these ~~These~~ include alternate timescales (clause 16.3 of [IEEE 1588-2008]), grandmaster clusters (clause 17.3 of [IEEE 1588-2008]), alternate master (clause 17.4 of [IEEE 1588-2008]), acceptable master table (clause 17.6 of [IEEE 1588-2008]), and the experimental cumulative frequency scale factor offset (Annex L of [IEEE 1588-2008]) all within [IEEE 1588].

For implementations compliant to [IEEE 1588-2019] these include unicast message negotiation (clause 16.1 of [IEEE 1588-2019]), alternate timescale offsets (clause 16.3 of [IEEE 1588-2019]), grandmaster clusters (clause 17.2 of [IEEE 1588-2019]), alternate master (clause 17.3 of [IEEE 1588-2019]), unicast discovery (clause 17.4 of [IEEE 1588-2019]), acceptable master table (clause 17.5 of [IEEE 1588-2019]), and the Cumulative frequency transfer method for synchronizing clocks (clause 16.10 of [IEEE 1588-2019]).

A.10   PTP common header flags

The PTP common header flag values, and whether or not each flag is used in this profile, are given in Table A.6.

NOTE – Some of these flags are used only in certain PTP messages, and not in all the PTP messages, see [IEEE 1588-2008] clause 13.3.2.6 or clause 13.3.2.8 of [IEEE 1588-2019].

For implementations compliant to [IEEE 1588-2008] the ~~The~~ following rule defined in [IEEE 1588-2008] clause 13.3.2.6, must be respected: "For message types where the bit is not defined in Table 20, the values shall be FALSE."

For implementations compliant to [IEEE 1588-2019] the following rule defined in clause 13.3.2.8 of [IEEE 1588-2019], must be respected: "For message types where the bit is not defined in Table 37 of [IEEE 1588-2019], the values shall be FALSE."

Table A.6 – PTP flags

Octet	Bit	Flag	Value to be sent	Behaviour for the receiving node
0	0	alternateMasterFlag	See clause 6.9 of this Recommendation	Used
0	1	twoStepFlag	As per PTP	Used
0	2	unicastFlag	TRUE	Used
0	5	PTP profile Specific1	FALSE	Flag is ignored
0	6	PTP profile Specific2	FALSE	Flag is ignored
0	7	Reserved	FALSE	Reserved by PTP and flag is ignored
1	0	leap61	As per PTP (Note 2)	Used
1	1	leap59	As per PTP (Note 2)	Used
1	2	currentUtcOffsetValid	As per PTP (Notes 2, 3)	Used (Note 4)

Table A.6 – PTP flags

1	3	ptpTimescale	TRUE	Used
1	4	timeTraceable	See Table 2	Used
1	5	frequencyTraceable	See Table 2	Used
1	6	(Note 1)	(Note 1)	(Note 1)

NOTE 1 – An additional flag "synchronizationUncertain" has been defined in Annex E; the use of the "synchronizationUncertain" flag is optional.

NOTE 2 – When a clock is in holdover, within holdover specification, the PTP clock may continue to advertise the last known leap second event. If there was no pending leap second event, then the PTP clock continues to advertise FALSE for the pending leap second fields (leap59 and leap61). If there was a pending leap second event, the PTP clock may choose to either advertise FALSE for the pending leap second fields immediately or continues to advertise the leap second event. In the latter instance the PTP clock would clear the leap59 and leap61 second event field(s) and adjust the UTC offset field at the appropriate time based on its local PTP time (i.e., the local PTP time's UTC timescale rolling over at UTC midnight) if the PTP clock is still in holdover, within holdover specification. When a clock is in holdover, out of holdover specifications, the PTP clock behaviour with respect to leap second event is implementation specific. It is recommended that PTP clock continue to advertise any upcoming leap second event as appropriate.

NOTE 3 – When a clock is in holdover, within holdover specification, the PTP clock may continue to advertise the last known UTC offset with UTC offset valid TRUE. If the last known UTC offset valid was FALSE, then the PTP clock continues to advertise FALSE. If the last known UTC offset valid was TRUE then the PTP clock may choose to either advertise UTC offset valid FALSE immediately (freezing the UTC offset value) or continues to advertise the last known UTC offset with UTC offset valid TRUE. The UTC offset field may be updated as described in Note 2 above. When a clock is in holdover, out of holdover specifications, the PTP clock behaviour with respect to UTC offset is implementation specific.

NOTE 4 – Usage of currentUtcOffset from an Announce message, which indicates currentUtcOffsetValid as FALSE, may lead to the wrong UTC time calculation.

Annex B

Options to establish the PTP topology with the Alternate BMCA

(This annex forms an integral part of this Recommendation.)

Options to establish the PTP topology with the Alternate BMCA are described in Annex C of [ITU-T G.8275].

Annex C

Inclusion of an external phase/time input interface on a PTP clock

(This annex forms an integral part of this Recommendation.)

The model for inclusion of a unidirectional, external phase/time interface on a PTP clock is provided in Annex B of [ITU-T G.8275].

Annex D

TLV for PTP interface rate (optional)

(This annex forms an integral part of this Recommendation.)

This annex is optional but, if implemented, it is necessary for the equipment to conform to requirements contained herein. When a PTP port in MASTER state that is providing timing service has a different interface rate than a PTP port in SLAVE state receiving the timing service, delay asymmetry may occur as described in [ITU-T G.8271] Appendix V 'Delay asymmetry resulting from interface rate change in PTP-unaware network elements'. If the slave clock is aware of both its own PTP port interface rate, as well as the master clock PTP port interface rate, then the slave clock may compensate for such delay asymmetry. The following TLV may be appended to a signalling message that contains GRANT_UNICAST_TRANSMISSION TLV so that the master clock may communicate its PTP port interface rate to the slave clock.

Table D.1 – INTERFACE_RATE TLV

Bits								Octets	TLV offset
7	6	5	4	3	2	1	0		
tlvType								2	0
lengthField								2	2
organizationId								3	4
organizationSubType								3	7
interfaceBitPeriod								8	10
numberBitsBeforeTimestamp								2	18
numberBitsAfterTimestamp								2	20

tlvType (Enum16)

The value of tlvType shall be the ORGANIZATION_EXTENSION value (0x0003)

lengthField (Uinteger16)

The value of lengthField shall be 18 bytes.

organizationId (Octet [3])

The value of organizationId shall be the OUI value assigned by ITU-T = 0x0019A7.

organizationSubType (Enum24)

The value of organizationSubType for the INTERFACE_RATE TLV shall be 0x000002.

interfaceBitPeriod (Uinteger64)

The period of 1-bit of the transmitting PTP timestamp interface, excluding line encoding. The value is encoded as an unsigned integer in units of attoseconds (10^{-18} s) to accommodate interface bit periods less than 1 ns.

numberBitsBeforeTimestamp (Uinteger16)

The length of the packet prior to the timestamp point, in bits.

numberBitsAfterTimestamp (Uinteger16)

The length of the packet after the timestamp point, in bits.

By way of example, the following values may be used for 1 GbE interface with Annex D encapsulation.

– tlvType = 0x0003

– lengthField = 18

– organizationId = 0x0019A7

– organizationSubType = 0x000002

– interfaceBitPeriod = 0x0000,0000, 3B9A, CA00

– numberBitsBeforeTimestamp = (8 bytes pre-amble x 8 bits/byte) = 64

– numberBitsAfterTimestamp = ((86 bytes payload + 4 bytes FCS) x 8 bits/byte) = 720

NOTE 1 – The supported interfaces (and interface speed) for an equipment clock are listed in the relevant equipment clock specification (which is for further study), and not in this profile.

NOTE 2 – The TLV and interfaceBitPeriod format is applicable to single-lane and mutli-lane interfaces.

Table D.2 shows information about various interface speeds and the appropriate interfaceBitPeriod value.

Table D.2 – Informational interface speeds and type mappings

Interface Speed	ns per bit	Atto-sec per bit	64-bit atto-sec Representation
1	1,000,000,000.000	10^{18}	0x0DE0,B6B3,A764,0000
10 M	100.000	100,000,000,000	0x0000,0017,4876,E800
100 M	10.000	10,000,000,000	0x0000,0002,540B,E400
1G	1.000	1,000,000,000	0x0000,0000,3B9A,CA00
10 G	0.100	100,000,000	0x0000,0000,05F5,E100
25G	0.040	40,000,000	0x0000,0000,0262,5A00
40G	0.025	25,000,000	0x0000,0000,017D,7840
100G	0.010	10,000,000	0x0000,0000,0098,9680
1 T	0.001	1,000,000	0x0000,0000,000F,4240

Annex E

Synchronization uncertain indication (optional)

(This annex forms an integral part of this Recommendation.)

The synchronization uncertain indication is described in Annex D of [ITU-T G.8275].

Annex F

Mapping from PTP clockClass values to quality levels

(This annex forms an integral part of this Recommendation.)

Referring to [ITU-T G.8275] Appendix IV 'Use cases for mapping from PTP clockClass values to quality levels' this annex covers the quality levels that a PTP clock will output on its frequency interface in various scenarios.

There are three scenarios:

- First, the output quality level when the PTP clock is synchronizing to an upstream PTP clock (i.e., the PTP clock parent data set is not that of the local clock) and the PTP clock does not have a physical layer frequency reference.

- Second, the output quality level when the PTP clock is synchronizing to an upstream PTP clock (i.e., the PTP clock parent data set is not that of the local clock) and the PTP clock does have a physical layer frequency reference.

- Third, the output quality level when the PTP clock is not synchronizing to an upstream PTP clock (i.e., the PTP clock parent data set is that of the local clock).

For the first scenario, Table F.1 maps the received clockClass value to an egress quality level.

Table F.1 – Mapping of clockClass values for the first scenario

PTP parentDS. grandmasterClockQuality. clockClass	PTP timePropertiesDS. frequencyTraceable flag from PTP parent	ITU-T G.781/G.8264 Option I QLs	ITU-T G.781/G.8264 Option II QLs
6	N/A	QL-PRC	QL-PRS
7	TRUE	QL-PRC	QL-PRS
7	FALSE	Note	Note
135	TRUE	QL-PRC	QL-PRS
135	FALSE	Note	Note
140	N/A	QL-PRC	QL-PRS
150	N/A	QL-SSU-A	QL-ST2
160	N/A	QL-SSU-B	QL-ST3E
165	N/A	QL-SEC/ QL-EEC1	QL-ST3/ QL-EEC2
248	N/A	QL-SEC/ QL-EEC1	QL-ST3/ QL-EEC2
255	N/A	QL-SEC/ QL-EEC1	QL-ST3/ QL-EEC2
NOTE – The PTP clock is synchronized to an upstream PTP clock that is non-traceable to a Category 1 frequency source, but is within holdover specification. It is implementation specific if the PTP clock sends QL PRC/PRS or some other value.			

The second scenario is for further study.

For the third scenario, Table F.2 maps the local clock's clockClass value to an egress quality level.

Table F.2 – Mapping of clockClass values for the third scenario

Phase/time traceability description	defaultDS. clockQuality. clockClass	ITU-T G.781/G.8264 Option I QLs	ITU-T G.781/G.8264 Option II QLs
T-GM connected to a PRTC in locked mode (e.g., PRTC traceable to GNSS)	6	QL-PRC	QL-PRS
T-GM in holdover, within holdover specification, traceable to Category 1 frequency source	7	QL-PRC	QL-PRS
T-GM in holdover, within holdover specification, non-traceable to Category 1 frequency source	7	Note	Note
T-BC-P/A in holdover, within holdover specification, traceable to Category 1 frequency source	135	QL-PRC	QL-PRS
T-BC-P/A in holdover, within holdover specification, non-traceable to Category 1 frequency source	135	Note	Note
T-GM in holdover, out of holdover specification, traceable to Category 1 frequency source	140	QL-PRC	QL-PRS
T-GM in holdover, out of holdover specification, traceable to Category 2 frequency source	150	QL-SSU-A	QL-ST2
T-GM in holdover, out of holdover specification, traceable to Category 3 frequency source	160	QL-SSU-B	QL-ST3E
T-BC-P/A in holdover, out of holdover specification	165	Note	Note
T-GM or T-BC-P/A without time reference since start-up	248	Note	Note
T-TSC-P/A, acting as an OC (does not send *Announce* messages)	255	Note	Note
NOTE – The egress QL is based on the category of the frequency traceable reference. If the clock is syntonized by a physical layer frequency source, then the egress QL is decided by the ingress QL of physical layer frequency source. If the clock is not syntonized (such as if it is relying solely on the local oscillator) then the quality of the local frequency clock is used.			

Appendix I

Considerations on the use of priority2

(This appendix does not form an integral part of this Recommendation.)

The PTP attribute priority2 is configurable in this profile. In some special circumstances, the use of the priority2 attribute can simplify the network management. This appendix describes two use cases; other possible cases are for further study.

Case 1

Operators can configure the PTP attribute priority2 to make all of the T-BC-Ps either traceable to one T-GM, or traceable to two different T-GMs at the same time.

G.8275.2-Y.1369.2(16)_FI.1
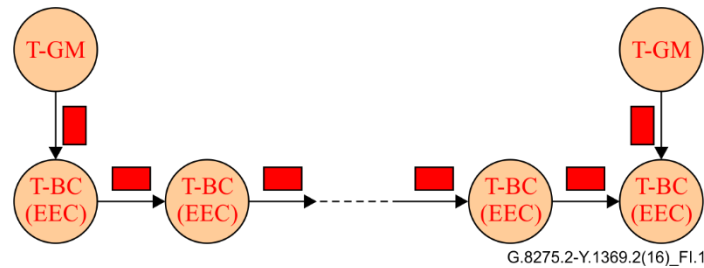
Figure I.1 – Use of priority2 with two T-GMs in the network

For example, in Figure I.1, if all other PTP attributes of the two T-GMs are the same, and the two T-GMs are configured with the same priority2 value, each T-BC-P will select the T-GM with the shortest path. If the two T-GMs are configured with different priority2 values, all of the T-BC-Ps will synchronize to the T-GM with the smallest priority2 value.

Case 2

Operators can configure the PTP attribute priority2 to prevent the T-BC-Ps of an upstream network from synchronizing with the T-BC-Ps of a downstream network when the T-GM is in failure.
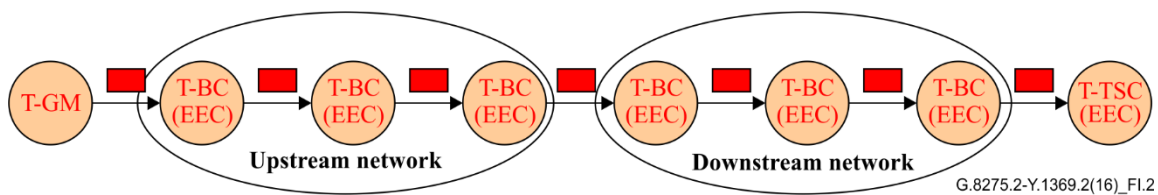

G.8275.2-Y.1369.2(16)_FI.2

Figure I.2 – Use of priority2 with T-BC-P/As of different network layers

For example, in Figure I.2, if all other PTP attributes of all of the T-BC-Ps are the same, and the PTP attribute priority2 of all of T-BC-Ps are configured with the same value, then when the T-GM is in failure, the T-BC-Ps in the upstream network can synchronize with the T-BC-Ps in the downstream network, depending on the clockIdentity values of all of the T-BC-Ps. If the T-BC-Ps in the upstream network are configured with a smaller priority2 value than the T-BC-Ps in the downstream network then, when the T-GM is in failure, the T-BC-Ps in the downstream network will synchronize to the T-BC-P s in the upstream network.

NOTE – The examples of this clause also apply when T-BC-As are deployed rather than T-BC-Ps.

Appendix II

Considerations on a T-TSC-A or T-TSC-P connected to an end application

(This appendix does not form an integral part of this Recommendation.)

The default T-TSC-A and T-TSC-P clockClass (248 for BC and 255 for OC) generally implies that the T-TSC-A or T-TSC-P will lock to the local time reference as a primary source of time (in case of APTS) or to an external PTP reference when available.

The actual synchronization source ultimately used by the end application depends on the applicable synchronization needs. This process is out of the scope of this recommendation.

As an example, the decision to use the PTP reference that has been selected by the T-TSC-A or T-TSC-P (e.g., instead of entering holdover), could depend on the actual clockQuality, frequencyTraceable flag and timeTraceable flag associated to the T-TSC-A or T-TSC-P input. Additional aspects as related to performance monitoring of the external reference might also be considered. This is implementation specific.

As an example, when it is required to meet the network timing requirements as per e.g., [ITU-T G.8271], it would be necessary that the external reference has clockClass 6, 7 or 135 and that the timeTraceable flag is TRUE in order to be used by the End Application. When this condition is not met, the end application may decide to enter holdover (either on the internal oscillator or driven by synchronous Ethernet).

NOTE – The specific behaviour for the T-TSC-P or T-TSC-A embedded in the end application is outside the scope of this Recommendation. It is assumed that interoperability with the profile is maintained.

Appendix III

PTP monitoring backup scenario example

*(This appendix does not form an integral part of this Recommendation.)*

Figure III.1 shows an example of a PTP deployment in steady-state, prior to a PTP connection failure between T-GM #A and T-BC-P #D. The following connectivity between the equipment is shown:

- T-GM #A is providing PTP service to T-BC-P #C
- T-GM #A is providing PTP service to T-BC-P #D
- T-BC-P #C is providing PTP service to T-TSC-P#E
- T-BC-P #D is providing PTP service to T-TSC-P #F

In addition, the figure shows two PTP connections that are not actively used for synchronization.

- T-BC-P #D is providing PTP service to T-BC-P #C, but T-BC-P #C is not selecting T-BC-P #D as the best PTP clock source. T-BC-P #C is providing PTP service to T-BC-P #D with the alternateMasterFlag set to TRUE on egress PTP messages to indicate its local port is in the PASSIVE state. T-BC-P #D is not selecting T-BC-P #C as the best PTP clock source.

The PTP connections between the T-BC-Ps allow the T-BC-Ps to support some specific types of monitoring. For example, the T-BC-P #D may monitor and learn the PDV characteristics of the PTP service from T-BC-P #C. This may be used to help the T-BC-P #D to synchronize more quickly to the T-BC-P #C backup PTP flow should the connection to the T-GM #A fail.
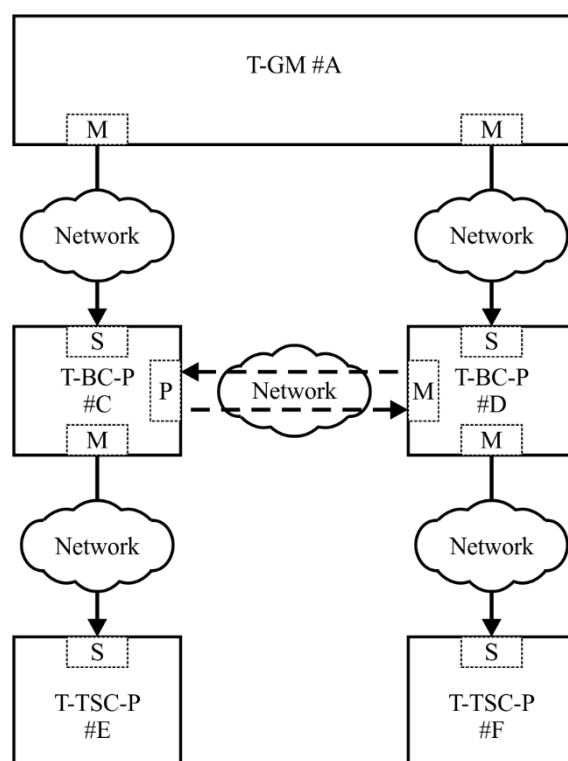
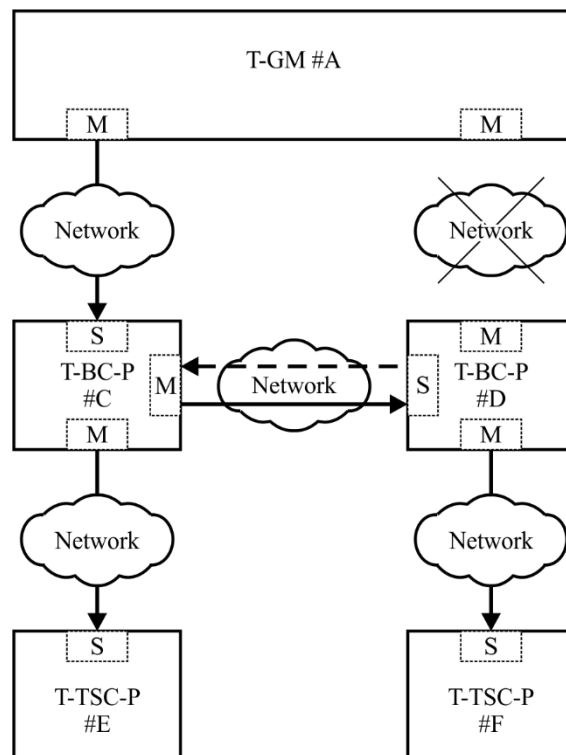G.8275.2-Y.1369.2(16)-Amd.1(17)_FIII.1

Figure III.1 − Steady-state, before A-D PTP connection failure

Figure III.2 shows the PTP deployment example in a steady-state, after the PTP connection failure between T-GM #A and T-BC-P #D. The T-BC-P #D takes advantage of the pre-failure monitoring of the T-BC-P #C PTP connection to enable a faster and less disruptive switching of the PTP service to an alternate source. After reaching steady-state again, the following connectivity takes place between the equipment:

• T-GM #A is providing PTP service to T-BC-P #C

• T-BC-P #C is providing PTP service to T-BC-P #D

• T-BC-P #C is providing PTP service to T-TSC-P#E

• T-BC-P #D is providing PTP service to T-TSC-P #F

In addition, there is one PTP connection that is not actively used for synchronization.

• T-BC-P #D is providing PTP service to T-BC-P #C with the alternateMasterFlag set to TRUE on egress PTP messages to indicate that its local port is in the SLAVE state. T-BC-P #C is not selecting T-BC-P #D as the best PTP clock source.

G.8275.2-Y.1369.2(16)-Amd.1(17)_FIII.2

Figure III.2 – Steady-state, after A-D PTP connection failure

In Figures III.1 and III.2 the inter-connection between T-BC-P #C and T-BC-P #D is shown using a single PTP port on T-BC-P #C and a single PTP port on T-BC-P #D. The inter-connection could alternatively be shown using two PTP ports; one PTP port on T-BC-P #C providing PTP service to one PTP port on T-BC-P #D, and a second PTP port on T-BC-P #D providing PTP service to a second PTP port on T-BC-P #C.

NOTE – The examples of this clause also apply when T-BC-As are deployed rather than T-BC-Ps.

Appendix IV

Description of PTP clock modes and associated contents of Announce messages

(This appendix does not form an integral part of this Recommendation.)

Description of PTP clock modes and associated contents of Announce messages are described in Appendix VIII of [ITU-T G.8275].

Appendix V

BMCA cycling between masters

(This appendix does not form an integral part of this Recommendation.)

This appendix describes a scenario where a PTP clock's BMCA may end up cycling (repeatedly switching) between master PTP clocks.

V.1 Scenario where a PTP clock's BMCA cycles between two masters

Consider the example scenario shown in Figure V.1 where a PTP clock has two candidate masters available in its unicast master table. In this example a T-GM #A has a better clockIdentity than a T-GM #B. Here only the Sync service is described (rather than both Sync & Delay_Resp service) to simplify the description.
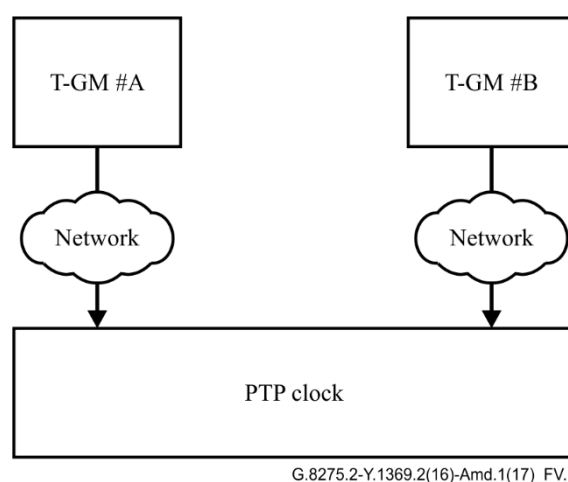

G.8275.2-Y.1369.2(16)-Amd.1(17)_FV.1

Figure V.1 – Example of a PTP clock with two candidate T-GMs

The following sequence of events may cause cycling of the PTP clock BMCA between a T-GM #A and a T-GM #B:

1) The PTP clock requests Announce service from a T-GM #A and a T-GM #B

2) The PTP clock selects a T-GM #A as a best master based on the Announce information content (clockIdentity in this example)

3) The PTP clock requests Sync service from a T-GM #A

4) The PTP clock does not get the Sync service from the T-GM #A, raising PTSF-lossSync
   • The PTSF-lossSync is a contributor to SF in the BMCA

5) The PTP clock selects T-GM #B based on T-GM-#A having PTSF alarm

6) The PTP clock requests Sync service from T-GM #B and receives that Sync service

7) The PTP clock chooses to cancel the Sync unicast session for T-GM #A as T-GM #A is no longer selected as best master

8) The PTP clock clears T-GM #A PTSF-lossSync alarm

9) The PTP clock selects T-GM #A based on better Announce clockClass **[repeats step #2]**.

V.2 Approaches to avoid a PTP clock's BMCA from cycling between two masters

This clause describes possible approaches to avoid or reduce the BMCA cycling phenomenon.

V.2.1 Requesting Sync and/or Delay_Resp service for non-selected masters

In order to avoid the cycling scenario, the PTP clock, when requesting Announce service from a non-selected master in the unicast master table, may request Sync and Delay_Resp service from that non-selected master. This allows constant monitoring of the non-selected masters' ability to deliver expected Sync and Delay_Resp service.

Additionally, the PTP clock, when de-selecting a master, may choose not to cancel the existing Sync and Delay_Resp unicast sessions.

As a result, when Sync service is not received from the non-selected master, the PTSF-lossSync would not be cleared when a master moves from selected to non-selected role.

With this approach, based on the above example, when the T-GM #A is not selected the PTP Clock will request (or not cancel) Sync service from T-GM #A. Given that the Sync service from T-GM #A is not received, the PTP clock will maintain the PTSF-lossSync alarm TRUE.

As a result the PTP clock will stay on T-GM #B until such time as T-GM #A Sync service is detected available (and PTSF-lossSync alarm is FALSE).

NOTE – This mechanism is also fully applicable to single Master (T-GM #A without T-GM #B) configurations. Is asserts PTSF-lossSync but allows detecting the return to normal behaviour by this Master.

V.2.2 Disqualify master triggering PTSF lossSync alarm

In order to avoid the cycling scenario, a slave can choose to cancel Sync and Delay_Resp unicast sessions with a master that has triggered a PTSF lossSync alarm. With this approach, the slave will not request Sync and Delay_Resp unicast sessions with any master in the unicast master table having a PTSF alarm set. Further, PTSF alarms on all masters will be cleared if and when the slave's local time source D0 is selected as best master by the BMCA (i.e., no external master is available) or by management. Any new master added to the unicast master table will start off with the PTSF flag set to FALSE, and the flag will remain FALSE until unicast negotiation is completed. As a result, any master that triggers a PTSF alarm will not be selected as best master until no other candidates are available. In order to use this feature, there must be a dedicated PTP port for each master in the unicast master table.

NOTE – This mechanism is not recommended for single Master (T-GM #A without T-GM #B) configurations. It would allow detecting the return to normal behavior by this Master, but also create endless and useless oscillations on the single Master's PTSF-lossSync.

Appendix VI

Considerations of PTP over IP transport in ring topologies

(This appendix does not form an integral part of this Recommendation.)

When using PTP messaging over an IP transport layer, there are some aspects of the Layer 3 protocol that need to be considered. The PTP layer delivers messages into the IP layer with a destination IP address. The IP layer then ensures the message is delivered to the destination as long as there is some path through the IP transport network from the source node to the destination address. The IP layer includes dynamic routing protocols that can adapt the path through the network based on available links between the IP routers. It can happen that the path taken by the IP transport layer may not be the path 'expected' by the synchronization planner. Applying some restrictions in the IP transport layer to control suboptimal paths for PTP messages may be beneficial. This is likely to be the case in ring topologies.

Taking the topology shown in Figure VI.1 as an example, the slave is configured to request unicast service from both BC3 and BC4. After receiving the Announce messages from both BC3 and BC4, the slave will run the BMCA and select BC4 as its parent clock based on the fact that the steps-removed value of BC4 is 1, compared to a steps-removed value of 3 for BC3. The slave would then request Sync messages from BC4.
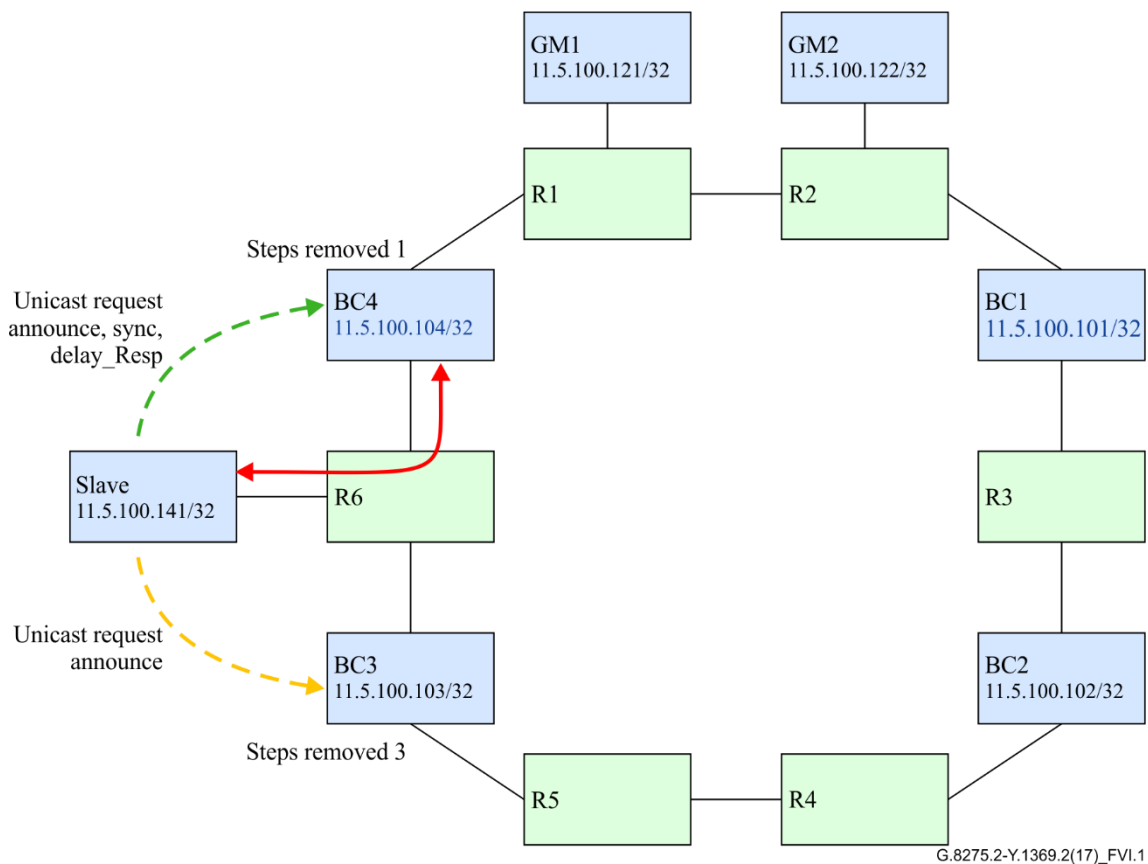

Figure VI.1 – Normal operation

If the connection between BC4 and R6 breaks (see Figure VI.2), then BC4 is not reached through the expected path. However, it can still be reached because routing protocols will retain the connection by routing the IP packets around the ring. BC4 is retained as the parent clock because it is still considered better by the BMCA.
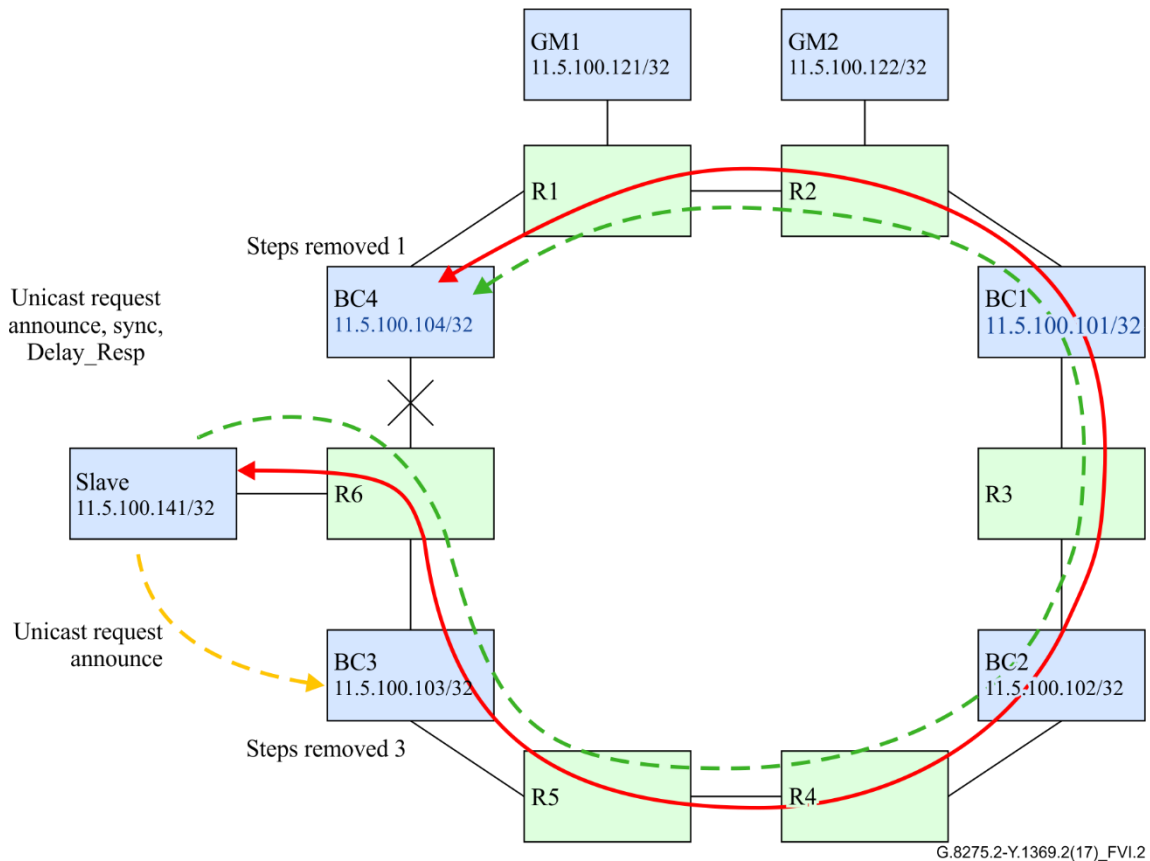
Figure VI.2 – Operation during a link failure between BC4 and R6

It is most likely that the desired operation is that the slave should switch to BC3 for better performance.

There are a few techniques that can be employed to ensure that in the failure scenario identified above, the slave will select BC3 as its parent clock. They are based on blocking the PTP IP messages from BC4 to the slave if those messages are transiting clockwise around the ring. The solution is based on blocking only the PTP messages and not the message of other protocols that might use the same IP addresses.

Option 1 – Unique IP addresses and static routes

In some deployment models, it may be possible to allocate unique IP addresses for the use of PTP alone. This then allows the use of static routes to control the direction of the PTP flows between the nodes. BC4 would be configured such that the only path to use to reach 11.5.100.141 (slave) would be the link between BC4 and R6. In addition, R6 could be configured such that the only path to use to reach 11.5.100.104(BC4) would be the link between R6 and BC4. If the link between R6 and BC4 fails, then there is no route available to get the IP packets between 11.5.100.141 and 11.5.100.104 so the slave will not receive Announces from BC4 and the BMCA will select BC3 as the parent clock. This is shown in Figure VI.3.
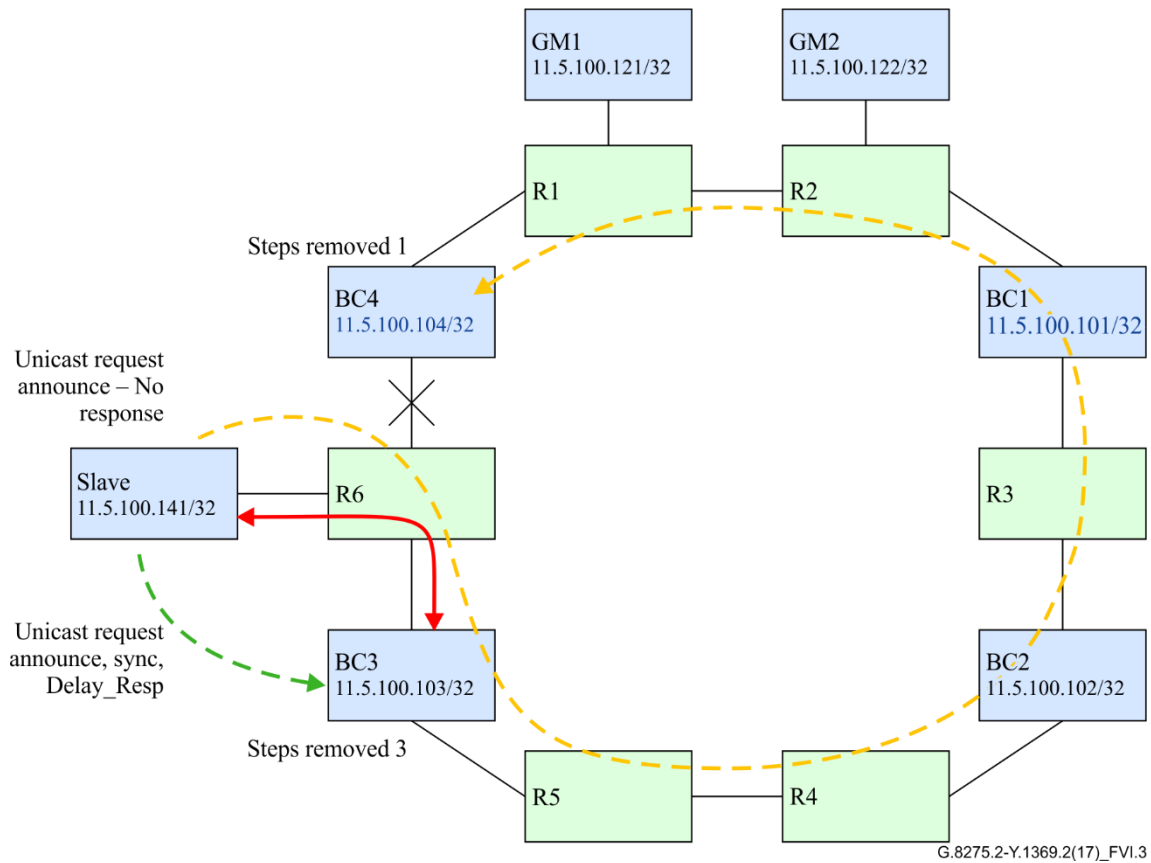
**Figure VI.3 – Operation during a link failure between BC4 and R6
when static routes are used**

Option 2 –IP filters

All routers support some level of IP filtering. Filters can be used to protect the control plane of the router from unwanted messages. They can be used in this case to control the acceptance of PTP messages on a subset of the routing interfaces.

In this case, R6 would be configured to protect the slave from PTP messages taking the wrong route. On the interface on R6 facing BC3, a filter could be applied to only allow messages to UDP port 319 or 320 if the source address matches that of the PTP process on BC3. Any messages sourced from BC4 that are received on that interface would be dropped. This is shown in Figures VI.4 and VI.5.
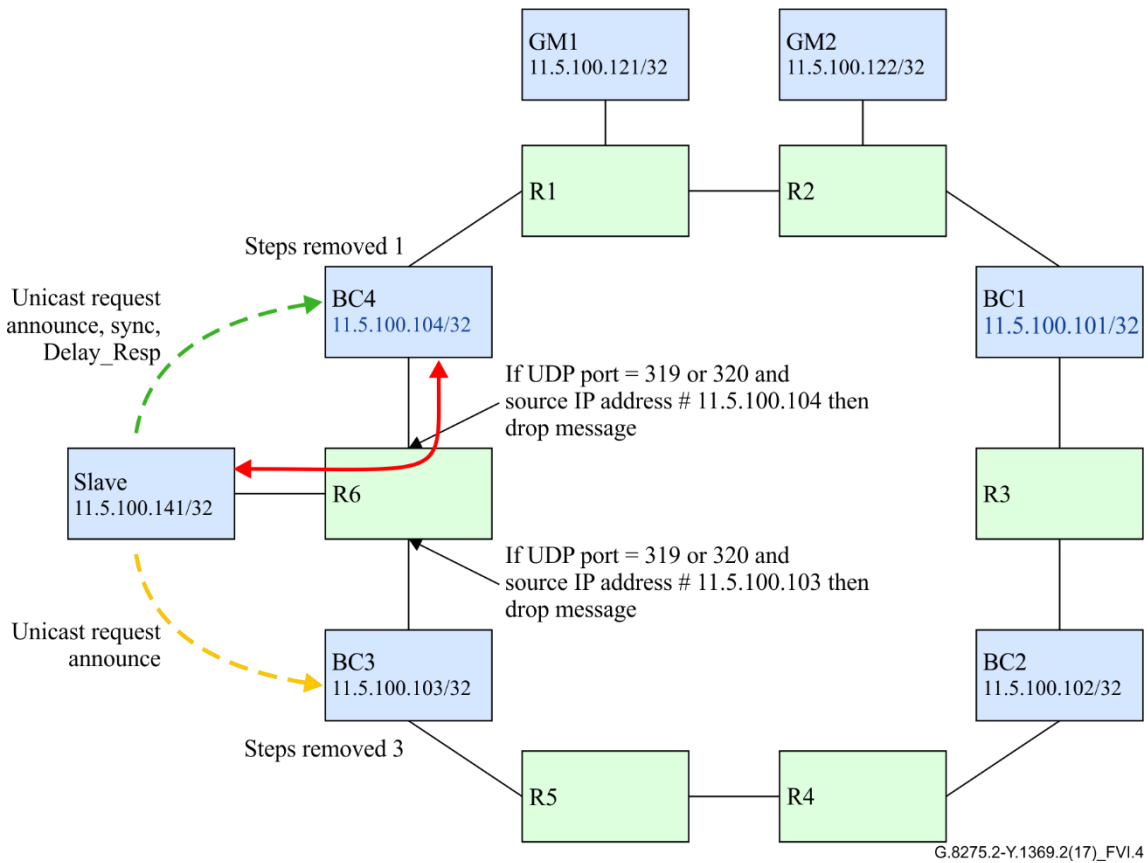
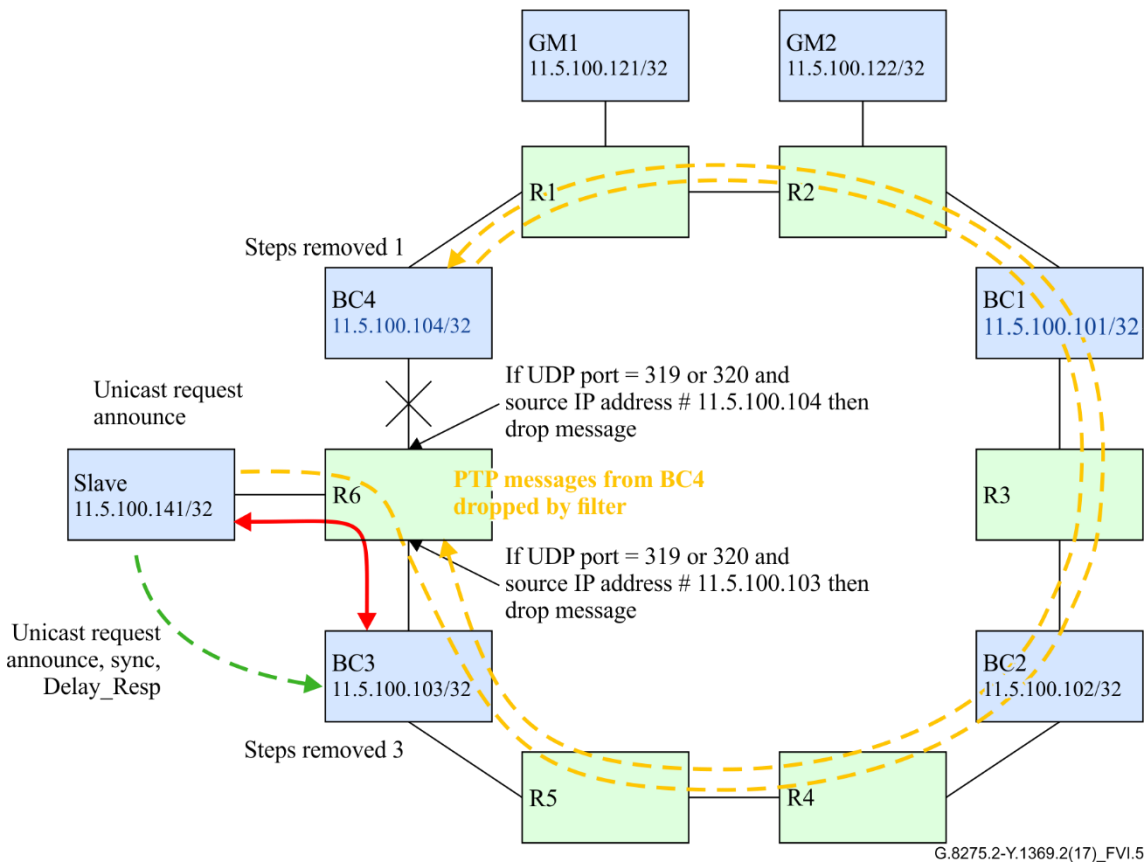**Figure VI.4 – Normal operation with IP filters in R6**

Figure VI.5 – Operation during a link failure between BC4 and R6 with IP filters in R6

Option 3 – BC processing of all PTP messages

A BC could terminate all PTP messages received into any of its ports for any domains used by the BC. Then the PTP messages could either be dropped or forwarded based on decisions within the PTP process itself. The choices might be to drop the message if the destination address of the PTP message was not an address owned by the BC or to deliver to the forwarding engine to be sent onward to the destination. The latter case might be used if the PTP message is for a different domain than the BC. Also in the latter case, the network element containing the BC might also update the correctionField of any forwarded event messages to compensate for the PTP message extraction and processing, i.e., support the transparent clock function for these messages. The message extraction from the IP plane can be accomplished if the router supports the policy based routing of IP packets.
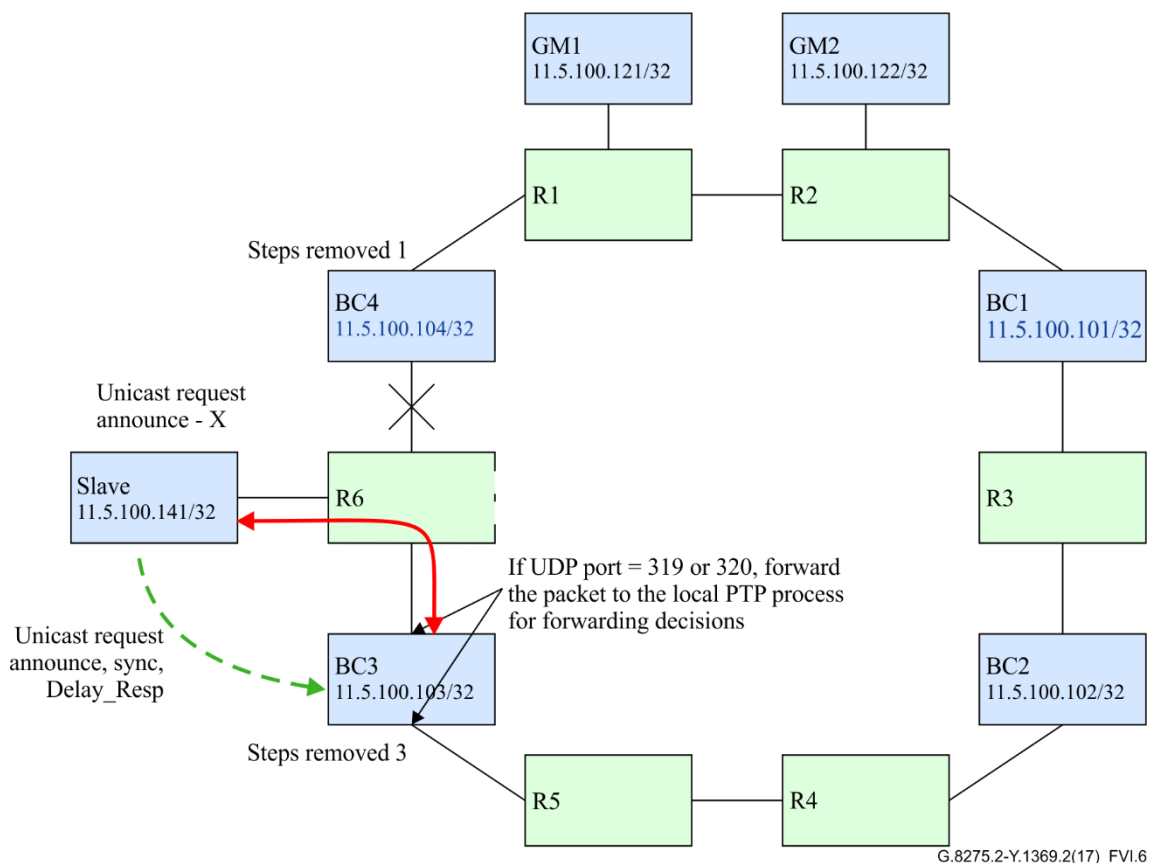
This example is shown in Figure VI.6.

Figure VI.6 – Operation during the link failure if all PTP messages are terminated

Option 4 – Use of the time to live (TTL) mechanism from IP transport

A PTP node may send PTP packets with the IP/Transport header carrying a time to live (TTL) field set to the minimum number of routing hops required to reach the peer PTP port with which it has a PTP contract. In a typical PTP-unaware network having unaware routers between master and slave, if the number of PTP unaware routers is larger than the TTL value of the PTP message, the PTP message will be dropped by one of the PTP-unaware routers. This can be used to limit the number of IP hops traversed by PTP packets between adjacent routers and avoiding communication through unwanted longer paths.

This behaviour may be per PTP port, or per PTP clock, and is implementation specific. It is assumed that in such a ring topology, IP routing will take care of ensuring that a shorter path to the PTP master is considered as a better route than the longer path around the ring.

As an example, if a slave clock has a directly connected master that can also be reachable through a longer path, it can use the TTL value of 1 to ensure that PTP packets reach the master only through the directly connected path rather than the longer path around the ring.

Appendix VII

Considerations on the configuration of PTSF-lossSync

(This appendix does not form an integral part of this Recommendation.)

The PTSF-lossSync may be set in the following scenarios. The operator should carefully consider which scenarios are relevant to the deployment:

– Failure to establish a contract with the best Master (as determined by the BMCA), due to no reply to a timing service request(s) or denial of a timing service request(s) (i.e., requests for timing messages, which include Sync, and Delay_Resp but exclude Announce).

  • To mitigate this issue, the implementation should generate the PTSF-lossSync alarm and select another available Master.

  • Note that setting a shorter timeout value for the reception of granted messages would allow for faster selection of another Master, limiting the holdover period.

– Total loss of timing service (Announce, Sync, and Delay_Resp) after a contract is established.

  • To mitigate this issue, the implementation should enter holdover and try selecting another available Master.

  • Note that setting a shorter timeout value for the reception of timing messages would allow for faster selection of another Master, limiting the holdover period.

  • The BMCA cycling issue, discussed in Annex V, proposes two possible ways to avoid endless alternating selection among Masters.

  • The PTSF-lossSync timer for missing timing messages (Sync and Delay_Resp) is independent from the one used for the Announce messages. Simultaneous loss of all messages may result in the triggering of PTSF-lossSync or announceReceiptTimeout, depending on which timer expires first.

As discussed in Appendix V, the PTSF-lossSync may be cleared after either normal delivery of timing service is restored, or other specific conditions to avoid the BMCA cycling issue occur.

Operators should carefully consider which scenarios are relevant to their deployments.

Appendix VIII

Operations over link aggregation

(This appendix does not form an integral part of this Recommendation.)

When two devices embedding PTP clocks compliant with this profile are connected via a link aggregation (LAG) as defined in [b-IEEE 802.1AX], each physical link should be accessed directly to transmit PTP messages, bypassing the LAG. This method prevents potential asymmetries that may be present when the forward and reverse paths are delivered over different links belonging to the LAG.

This functional model is only a suggested approach to implement this behaviour.

VIII.1 Functional model

One way of selecting the same link under a LAG bundle between two nodes R1 and R2 (see Figure VIII.1), is to explicitly map/configure the child interfaces/links under a given LAG bundle as primary, secondary and/or tertiary interfaces for the exchange of PTP packets.

By doing this, both ends of LAG bundle (R1 and R2) will deterministically select the same link for the PTP packet exchange (both in Tx and Rx direction). With the order of configuration, on failure of primary link, the secondary link will be selected, and on failure of secondary link, the tertiary link will be selected.

NOTE 1 – This functional model considers only two T-BC-Ps or T-BC-P and T-TSC-P connected over direct links using a LAG bundle, there is no PTP unaware nodes between T-BC-Ps or T-BC-P and T-TSC-P.

NOTE 2 – It is optional to specify more than two child interfaces for exchange of PTP packets in a LAG bundle.

NOTE 3 – Specifying primary and secondary links is only applicable for the exchange of PTP packets. It will not alter any data or affect other non-PTP packets transmitted over the LAG bundle nor will it disturb any link selection algorithms used for other protocols.
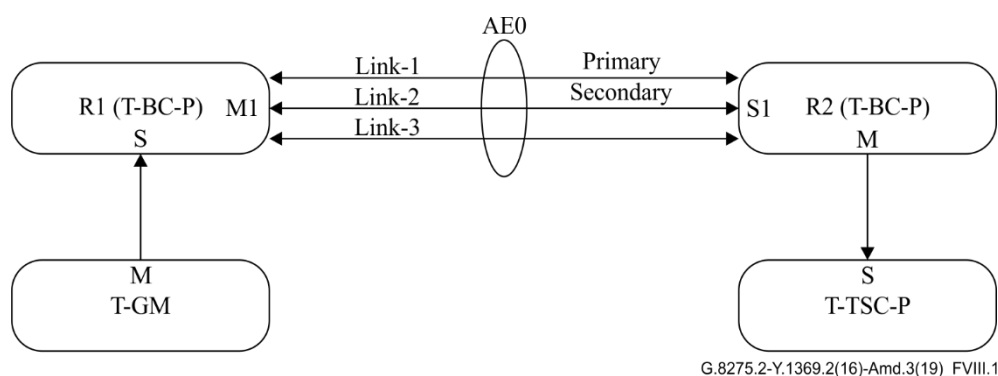
G.8275.2-Y.1369.2(16)-Amd.3(19)_FVIII.1

Figure VIII.1 – LAG bundle link selection example

In both the R1 and R2 nodes, the PTP port is configured on the AE0 bundle port, with child Link-1 as the primary path for PTP packets and child Link-2 as the secondary path. While the primary link (child Link-1) is active, both R1 and R2 will exchange all event and non-event PTP packets over this link. When the primary link (Link-1) fails, R1 and R2 would detect the link loss and both nodes would failover to use the configured secondary link (Link-2) for the transmission of PTP packets within the LAG bundle.

There are two cases when failed primary link recovers back:

VIII.2   Scenario

Assume primary Link-1 failed, R1 and R2 exchanges PTP packets over secondary link (Link-2). Now primary link (Link-1) comes back up.

Case-1: Revertive

If revertive mode is configured, when the failed Primary link (Link-1 in this example) comes back up, both R1 and R2 will switch back from the secondary link to the primary link for the exchange of PTP packets. This will be the default behaviour unless explicitly configured as non-revertible (case 2).

Case-2: Non-revertive

If non-revertive mode is configured, then when the failed primary link comes back up, the R1 and R2 nodes will continue to exchange PTP packets over the secondary link without reverting to the primary link.

NOTE 4 – In case of PTP over IP over a Layer 3 LAG bundle, it is not possible to configure PTP directly on the child interfaces of the LAG bundle, as child interfaces are L2 interfaces. It is necessary to enable PTP on the AE (Aggregated Ethernet) or bundled interface and further specify which child interfaces (i.e., primary, secondary, …) should be used to exchange the PTP packets over the LAG bundle. When configured correctly at both ends, the same link would be used to exchange PTP packets in the forward and reverse direction that will help to minimize the link asymmetry.

Appendix IX

Considerations on the use of [IEEE 1588-2019]

(This appendix does not form an integral part of this Recommendation.)

Considerations on the use of [IEEE 1588-2019] are discussed in Appendix IX of [ITU-T G.8275].

Appendix X

Considerations on selecting time out values

(This appendix does not form an integral part of this Recommendation.)

Careful consideration should be used when configuring syncReceiptTimeout and delayReqReceiptTimeout in a partially aware deployment. In a partially aware deployment there is a higher probability of loss of packets (or a burst of loss of packets) compared with fully aware networks. This leads to the desire to choose a good value for how many successive packet losses a PTP clock may tolerate before declaration of a receipt timeout, which may lead to raising a PTSF alarm, which in turn may lead to disqualification of the PTP connection from consideration for selection by the BMCA.

Typically, the length of the time a PTP clock can tolerate loss of Sync and Delay_Resp messages is related to the target performance requirements and also the design implementation of the PTP clock (such as the stability of the PTP clock's local oscillator to maintain the performance in-between the reception of PTP synchronization messages). The performance requirements and design implementation is known prior to PTP connection establishment.

The PTP clock receipt timeout properties are therefore configured by considering two additional related parameters

- Negotiated PTP message rate
- Duration of PTP message loss before entering a holdover state (whether holdover in-spec or holdover out-of-spec)

Note that the PTP clock that acts as the receiver of the PTP information is the one that initiates the contract negotiation with a request for a specific PTP message rate (per message type). With that negotiated PTP message information, combined with the known performance target and equipment design, it is then possible to determine the receipt timeout property for the connection.

X.1   Example receipt timeout calculation (Dynamic message rate)

A PTP clock may use the following formulas for setting values on the portDS.syncReciptTimeout and portDS.delayRespReceiptTimeout data set members to achieve desired performance, based on the message rate granted by unicast negotiation. Here, the Tolerated Consecutive Synchronization Message Loss Time (TLT) must be known in advance, based on the performance target and the equipment design. The TLT is defined as the time period the PTP clock can be without Sync and/or Delay_Resp messages and not enter the holdover state (whether holdover in-spec or holdover out-of-spec). In addition, for a proper error report management, there should also be a reasonable upper time limit to the receipt timeout; as described in Appendix X.3 this is assumed to be 15 minutes. Using the 15 minutes as a reasonable time limit to the receipt timeout, the maximum TLT should be 900 seconds.

TLT = Tolerated Consecutive Synchronization Message Loss Time [s]

NRS = Negotiated Message Rate PTP Sync messages [messages/s]

NRD = Negotiated Message Rate PTP Delay_Resp messages [messages/s]

portDS.syncReceiptTimeout = The minimum of {TLT × NRS, 65535}

portDS.delayRespReceiptTimeout = The minimum of {TLT × NRD, 65535}

Example ~**5** minutes TLT and **64** messages/s:

portDS.syncReceiptTimeout = The minimum of {5×60×64, 65535}

portDS.syncReceiptTimeout = 19200

Example ~**100** seconds TLT and **16** messages/s:

portDS.syncReceiptTimeout = The minimum of {100×16, 65535}

portDS.syncReceiptTimeout = 1600

X.2   Example receipt timeout calculation (Fixed message rate)

PTP clocks that operate only at one specific PTP message rate may use fixed values of the portDS.syncReciptTimeout and portDS.delayRespReceiptTimeout data set members that are implementation specific.

portDS.syncReceiptTimeout = fixed value by design

portDS.delayRespReceiptTimeout = fixed value by design

X.3   Range of receipt timeout

In some deployments that reference [b-ITU-T G.7710] on common equipment management function requirements, it may be desirable to report on loss of Sync or Delay_Resp messages within 15 minutes of the failure event. This may be a reasonable upper limit to the receipt timeout configurable range in a partially aware network deployment. When operating at 128 messages per second, with a receipt timeout of 65535, yields about 8.5 minutes. When operating at lower message rates, care should be taken not to set the receipt timeout to the maximum allowed within the receipt timeout range to avoid unreasonably large timeout values in units of seconds.

Bibliography

[b-ITU-T G.7710]   Recommendation ITU-T G.7710/Y.1701 (2020), *Common equipment management function requirements.*

[b-IEEE 802.1AX]   IEEE 802.1AX (2014), *IEEE Standard for Local and metropolitan area networks – Link Aggregation*.

[b-IEEE 1588-2019]   IEEE 1588 (2019), *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems.*

ITU-T Y-SERIES RECOMMENDATIONS

GLOBAL INFORMATION INFRASTRUCTURE, INTERNET PROTOCOL ASPECTS, NEXT-GENERATION NETWORKS, INTERNET OF THINGS AND SMART CITIES

GLOBAL INFORMATION INFRASTRUCTURE	
General	Y.100–Y.199
Services, applications and middleware	Y.200–Y.299
Network aspects	Y.300–Y.399
Interfaces and protocols	Y.400–Y.499
Numbering, addressing and naming	Y.500–Y.599
Operation, administration and maintenance	Y.600–Y.699
Security	Y.700–Y.799
Performances	Y.800–Y.899
INTERNET PROTOCOL ASPECTS	
General	Y.1000–Y.1099
Services and applications	Y.1100–Y.1199
Architecture, access, network capabilities and resource management	Y.1200–Y.1299
Transport	**Y.1300–Y.1399**
Interworking	Y.1400–Y.1499
Quality of service and network performance	Y.1500–Y.1599
Signalling	Y.1600–Y.1699
Operation, administration and maintenance	Y.1700–Y.1799
Charging	Y.1800–Y.1899
IPTV over NGN	Y.1900–Y.1999
NEXT GENERATION NETWORKS	
Frameworks and functional architecture models	Y.2000–Y.2099
Quality of Service and performance	Y.2100–Y.2199
Service aspects: Service capabilities and service architecture	Y.2200–Y.2249
Service aspects: Interoperability of services and networks in NGN	Y.2250–Y.2299
Enhancements to NGN	Y.2300–Y.2399
Network management	Y.2400–Y.2499
Computing power networks	Y.2500–Y.2599
Packet-based Networks	Y.2600–Y.2699
Security	Y.2700–Y.2799
Generalized mobility	Y.2800–Y.2899
Carrier grade open environment	Y.2900–Y.2999
FUTURE NETWORKS	Y.3000–Y.3499
CLOUD COMPUTING	Y.3500–Y.3599
BIG DATA	Y.3600–Y.3799
QUANTUM KEY DISTRIBUTION NETWORKS	Y.3800–Y.3999
INTERNET OF THINGS AND SMART CITIES AND COMMUNITIES	
General	Y.4000–Y.4049
Definitions and terminologies	Y.4050–Y.4099
Requirements and use cases	Y.4100–Y.4249
Infrastructure, connectivity and networks	Y.4250–Y.4399
Frameworks, architectures and protocols	Y.4400–Y.4549
Services, applications, computation and data processing	Y.4550–Y.4699
Management, control and performance	Y.4700–Y.4799
Identification and security	Y.4800–Y.4899
Evaluation and assessment	Y.4900–Y.4999

For further details, please refer to the list of ITU-T Recommendations.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	Tariff and accounting principles and international telecommunication/ICT economic and policy issues
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	**Transmission systems and media, digital systems and networks**
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Environment and ICTs, climate change, e-waste, energy efficiency; construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling, and associated measurements and tests
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	**Global information infrastructure, Internet protocol aspects, next-generation networks, Internet of Things and smart cities**
Series Z	Languages and general software aspects for telecommunication systems